

Altiris™ Monitor Solution for Servers 7.1 SP2 from Symantec™ User Guide



Altiris™ Monitor Solution for Servers 7.1 SP2 from Symantec™ User Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

See "[Altiris™ Monitor Solution for Servers 7.1 SP2 Symantec™ Third-Party Legal Notices](#)" on page 163.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our Web site at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan

customercare_apac@symantec.com

Europe, Middle-East, and Africa

semea@symantec.com

North America and Latin America

supportsolutions@symantec.com

Contents

Technical Support	4
Chapter 1 Introducing Monitor Solution	13
About Monitor Solution	13
What's new in Monitor Solution 7.1 SP2	14
Components of Monitor Solution	14
How Monitor Solution works	16
About Monitor Pack for Servers	17
Where to get more information	19
Chapter 2 Configuring the Monitor Solution Server	21
Configuring the monitor server	21
About monitor packs	22
Importing monitor packs	22
About database maintenance	23
Maintaining collected performance data	25
About heartbeat	26
Setting up the monitor server's heartbeat settings	27
Chapter 3 Configuring the Monitor Plug-in	29
About the Monitor Plug-in	30
About Monitor Plug-in installation policies	30
About Monitor Plug-in configuration policies	31
About Monitor policies	31
About Monitor Plug-in profiling	32
Preparing managed computers for agent-based monitoring	33
Creating new Monitor Plug-in settings	34
Configuring Monitor Plug-in settings	35
Monitor Plug-in configuration settings: General tab	36
Monitor Plug-in configuration settings: Performance Tuning tab	38
Monitor Plug-in configuration settings: Data Collection tab	39
Monitor Plug-in configuration settings: Maintenance Windows tab	42

Chapter 4	Configuring agentless monitoring	47
	About agentless monitoring	47
	About agentless monitoring and network discovery	48
	About monitor service	48
	Setting up a remote monitoring site server	49
	Installing the Pluggable Protocols Architecture (PPA) client computer component on a site server	51
	Removing monitor service from a site server	52
	Adding monitor service to a site server	53
	Monitor site server reports	54
	Configuring remote monitoring server settings	55
	Remote Monitoring Server Settings: General tab	55
	Remote Monitoring Server Settings: Performance Tuning tab	57
	Remote Monitoring Server Settings: Data Collection tab	59
Chapter 5	Working with Monitor Policies	63
	Creating a monitor policy with the monitor policy wizard	63
	Creating a monitor policy	65
	Editing agent-based monitor policies	66
	Editing agentless monitor policies	67
	Adding rules to a monitor policy	68
	About application detection	68
	Adding application detection to a monitor policy	69
	Application detection types	70
	How operators work in application detection	72
	Adding computers to a monitor policy	72
Chapter 6	Working with Metrics	75
	About metrics	75
	Creating and editing metrics in the metric library	75
	Adding a metric to a rule	77
	Types of metrics	77
	New COM Metric page	80
	New Command Metric page	80
	New Compound Metric page	84
	New Custom DLL Metric page	85

New Group Metric page	86
New HTTP Metric page	87
New Log Event Metric page	89
New Ping Metric page	91
New Performance Counter Metric page	92
New Port Metric page	93
New SNMP Metric page	94
New Smart Metric page	96
New SQL Metric page	96
New Windows Process Metric page	98
New Windows Service Metric page	99
New WMI Metric page	100
New WS-MAN Metric page	102
About multiple instance metrics	104
Chapter 7	
Working with rules	107
About rules	107
Creating and editing rules in the rule library	108
Cloning rules	110
Creating and editing rules	110
About metric evaluation	111
New Metric Evaluation page for log event rules	112
New Metric Evaluation page for Metric rule types	113
New Metric Evaluation page for NT Event rule types	114
Types of rules	115
Chapter 8	
Working with tasks and actions	117
About Monitor Solution tasks and actions	117
About severity states	119
Monitor task types	120
Adding tokens to a Send Email task	123
Adding actions to rules	124
Adding actions to monitor policies	125
Monitor client and server token types	127
Chapter 9	
Viewing Monitored Data	131
About the Monitoring and Alerting home page	131
Viewing historical performance data	132
Viewing real-time performance data	133
Monitored Resources dialog box and Resources with Historical Data dialog box	134

	Viewing Monitor Solution reports	134
Chapter 10	Using alert management	135
	About alerts	136
	About alert management	137
	Note on time zones and alerts	137
	About Event Console alert filters	138
	Alert Filter Settings page	140
	Filtering alerts	140
	Creating and saving alert filters	141
	About advanced search filters	142
	Creating advanced search filters	144
	Viewing alerts	145
	Hiding resolved alerts	146
	Alert Rule Settings page	146
	Creating an alert matching rule	147
	Adding or editing rules to discard alerts	148
	Forwarding alerts to another management system	149
	Running a task in response to an alert	149
	About Event Console tokens	150
	Event Console token types	150
	About the Event Console workflow rule	151
	About workflow rule configuration	152
	Adding or editing workflow rules	153
	About alert purging	154
	Purging old and low-severity alerts	155
	Viewing the health of an organizational group	155
	Working with Event Console tasks	156
	Change alert status task page	157
	Create resource task page	157
	Event Console purge policy task page	157
	Raise message task page	157
	Reprioritize alert task page	158
Index		159
Appendix A	Altiris™ Monitor Solution for Servers 7.1 SP2 Symantec™ Third-Party Legal Notices	163
	Third-Party Legal Attributions	163
	Expat XML Parser v2.0.1	163
	Net-SNMP v 5.4.1	164
	RegExp	170

RegExp License	170
----------------------	-----

Introducing Monitor Solution

This chapter includes the following topics:

- [About Monitor Solution](#)
- [What's new in Monitor Solution 7.1 SP2](#)
- [Components of Monitor Solution](#)
- [How Monitor Solution works](#)
- [About Monitor Pack for Servers](#)
- [Where to get more information](#)

About Monitor Solution

Monitor Solution lets you monitor various aspects of computer operating systems, applications, and devices. These aspects can include events, processes, and performance. This ability helps you ensure that your servers and your devices work and reduces the costs of server and network monitoring.

Monitor Solution lets you do the following tasks:

- Identify the health of your environment by collecting detailed data from servers, applications, and network devices.
- Analyze trends and isolate recurring issues by collecting comprehensive real-time and historical performance data.
- Pinpoint problems, define their cause, and take automated actions to resolve them.

Monitor Solution supports both agent-based and agentless monitoring methods. It runs on the Symantec Management Platform and is a key component of Server Management Suite.

See “[Components of Monitor Solution](#)” on page 14.

See “[How Monitor Solution works](#)” on page 16.

What's new in Monitor Solution 7.1 SP2

In the 7.1 SP2 release of Monitor Solution, the following new features are introduced:

Table 1-1 List of new features

Feature	Description
Support for new platforms.	<p>You can install Monitor Agents on the computers that are running the following platforms:</p> <ul style="list-style-type: none"> ■ Red Hat Enterprise Linux 6.0 x86/x64 ■ Red Hat Enterprise Linux 6.1 x86/x64 ■ SUSE Linux Enterprise Server 11 SP1 x86/x64 ■ Solaris 10 Update 7
Changes on the Monitoring and Alerting Portal page.	<p>The following changes were made on the Monitoring and Alerting Portal page:</p> <ul style="list-style-type: none"> ■ The Activated Monitor Policies Web part was removed. However, it is still available under Settings > Notification Server > Console Settings > Web Parts > Monitoring. ■ Monitor Site Servers Status Web part was added. ■ Group View - Aggregate health by resource Web part was added.
A new report.	<p>The Server Health Report Summary report was added. The report is available under Reports > Monitoring and Alerting > Monitor > Servers > Common Performance > System.</p>

See “[About Monitor Solution](#)” on page 13.

Components of Monitor Solution

Monitor Solution lets you monitor different aspects of servers and applications. This monitoring is done through multiple monitoring solutions that work together using a common set of Monitor Solution components that are called the core components. Each monitoring solution uses the core components and includes a

set of monitoring components specific to the purpose of the monitoring solution. Each solution also includes numerous reports to help analyze data. This separation of core functionality provides flexibility to comprehensively monitor aspects of computer resources and network devices.

The core components of Monitor Solution are as follows:

- **Monitor Plug-in**

The Monitor Plug-in performs the monitoring work on computers. The Monitor Plug-in is a plug-in to the Symantec Management Agent, which is installed on monitored computers. The Monitor Plug-in receives configuration data from the Notification Server computer specifying what aspects of the computer are to be monitored.

See “[About the Monitor Plug-in](#)” on page 30.

- **Agentless monitoring**

A monitor service on a site server acts in place of a Monitor Plug-in. It lets you monitor certain aspects of your computers that cannot have plug-ins installed on them.

See “[About agentless monitoring](#)” on page 47.

See “[About monitor service](#)” on page 48.

- **Real-time and historical performance viewers**

Performance monitoring lets you view the performance of a computer in real time or historically. This data makes it easy to analyze performance and identify problems.

See “[Viewing real-time performance data](#)” on page 133.

See “[Viewing historical performance data](#)” on page 132.

- **Reports**

Numerous predefined reports help you analyze your data; you can also create custom reports if the predefined reports do not meet your needs.

See “[Viewing Monitor Solution reports](#)” on page 134.

- **Monitor packs**

Monitor packs include the necessary monitor policies, metrics, rules, and tasks for monitoring an operating system or application. Monitor packs also contain preconfigured monitor policies with preset thresholds and severities.

See “[About monitor packs](#)” on page 22.

- **Monitor policy**

A monitor policy is group of monitoring rules. You apply monitor policies to the groups of computers and devices that you want to monitor. Monitor policies inform the Monitor Plug-in or the Remote Monitoring Server of what data you want monitored and how that data should be analyzed. The data is evaluated against the conditions of rules. Based on these rules the Monitor Plug-in can

run automated actions in response to data that reaches an undesired state or range. The Monitor Plug-in returns the monitored data to the Notification Server computer. The Notification Server computer uses monitored data to run Task Server tasks for real-time performance monitoring and historical performance reporting.

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

- Rules

Rules specify how to analyze the metric data or the event data that the Monitor Plug-in and the Remote Monitoring Server collect. Rules also define under what conditions they are triggered and the actions taken.

See “[About rules](#)” on page 107.

- Actions and Tasks

You can add actions and tasks to a rule or a policy. Rules are triggered when monitored metric data reaches a determined value or goes beyond an acceptable value range. A triggered rule sends an alert, and any actions or tasks that are specified for that rule or policy are executed. Monitor actions and tasks can also be scheduled or run on demand. You can run tasks from a task server or you can choose from several Monitor Plug-in-specific task types.

See “[About Monitor Solution tasks and actions](#)” on page 117.

- Metrics

Metrics define how a Monitor Plug-in or the Remote Monitoring Server collects data from supported data sources, called metric sources. Each plug-in can use numerous metrics to define all of the data that you want to collect.

See “[About metrics](#)” on page 75.

See “[About Monitor Solution](#)” on page 13.

See “[How Monitor Solution works](#)” on page 16.

How Monitor Solution works

Monitor Solution continuously collects and analyzes data that is captured from computers and other devices on your network. When data is captured that meets a criteria that you predefine, alerts can be raised to notify you and actions can be taken.

The Monitor Plug-in or the Remote Monitoring Server gathers the data that you want to monitor. The data is remotely managed from the Symantec Management Console. The Monitor Plug-in and the Remote Monitoring Server receive instructions from the Notification Server computer. These instructions are called monitor policies. Monitor policies instruct the plug-in and Remote Monitoring

Server of what and how you want to monitor, and what actions you want to be done.

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

Monitor policies are built from metrics and rules. Metrics specify the data to collect, and rules specify when you want to be notified and what you want to do.

See “[About metrics](#)” on page 75.

See “[About rules](#)” on page 107.

Multiple monitor policies with similar purposes, such as monitoring an operating system or an application, are grouped together into monitor packs. Monitor packs must be imported into the Notification Server computer following the installation of Monitor Solution.

See “[About monitor packs](#)” on page 22.

Monitor Solution uses triggered rules to take automated actions called tasks.

See “[About Monitor Solution tasks and actions](#)” on page 117.

Tasks can include sending an email, creating an alert in the Event Console, or running a command on the monitored computer. The Monitor Plug-in or the Notification Server computer can execute tasks.

See “[About Monitor Solution](#)” on page 13.

See “[Components of Monitor Solution](#)” on page 14.

About Monitor Pack for Servers

The Monitor Pack for Servers component provides a collection of monitor packs that monitors the health of your servers. Monitor packs are the containers of the monitor policies that monitor services and events of the server health, operating system, and applications.

The Monitor Pack for Servers component contains both agent-based and agentless monitoring policies. You can run agent-based monitor policies on the computers that have the Monitor Plug-in installed. Agentless monitor policies let you monitor resources without the Monitor Plug-in.

See “[Preparing managed computers for agent-based monitoring](#)” on page 33.

See “[Configuring the monitor server](#)” on page 21.

Depending on the aspects that you want to monitor, you can enable or disable the policies that are included in the monitor packs. You can also create new policies. Each monitor policy contains the necessary rules, metrics, and tasks that let you monitor your resources. Rules and metrics let you define the metric evaluation

and metric data that you want to monitor. Tasks let you specify the automated actions that occur when the metric data reaches certain evaluation.

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

See “[About metrics](#)” on page 75.

See “[About rules](#)” on page 107.

See “[About Monitor Solution tasks and actions](#)” on page 117.

The Monitor Pack for Servers component also includes numerous reports that help you analyze the data and tune the performance of your servers.

See “[About the Monitoring and Alerting home page](#)” on page 131.

Table 1-2 Default monitor packs included in the Monitor Pack for Servers component

Monitor pack	Description
AIX - Basic	This monitor pack lets you monitor disk, memory, network, processor, and other aspects of AIX servers.
ESX - Basic	This monitor pack lets you monitor the health and performance of ESX servers including disk, memory, network, and processor.
ESX - Extended Host	This monitor pack lets you monitor the ESX host servers for virtualization metrics including host disk, virtual memory, system and virtual processor.
Linux - Basic	This monitor pack lets you monitor disk, memory, network, processor, and other aspects of Linux servers.
Linux Server Health	This monitor pack lets you monitor health and performance of your Linux Servers. This pack is a single policy that you can apply to all your Linux Servers to quickly apprehend the operation system health and performance.
Solaris - Basic	This monitor pack lets you monitor disk, memory, network, processor, and other aspects of Solaris servers.
Windows 2003	This monitor pack lets you monitor the health and performance on the Windows 2003 servers including disk, memory, network, and processor.
Windows 2008	This monitor pack lets you monitor the health and performance on the Windows 2008 servers including disk, memory, network, and processor.

Table 1-2 Default monitor packs included in the Monitor Pack for Servers component (*continued*)

Monitor pack	Description
Windows Agentless Policy	This agentless monitor pack lets you monitor the availability and performance on the Windows 2003/2008 servers including disk, memory, network, and processor. The agentless monitor policy lets you monitor computers without the Monitor Plug-in installed. Because the Monitor Plug-in is not available, fewer aspects of the computers are available to be monitored.
Windows Server Health	This monitor pack lets you monitor health and performance of your Linux Servers. This pack is a single policy that you can apply to all your Linux Servers to quickly apprehend the operation system health and performance.

Where to get more information

Use the following documentation resources to learn about and use this product.

Table 1-3 Documentation resources

Document	Description	Location
Release Notes	Information about new features and important issues.	The Supported Products A-Z page, which is available at the following URL: http://www.symantec.com/business/support/index?page=products Open your product's support page, and then under Common Topics , click Release Notes .
User Guide	Information about how to use this product, including detailed technical information and instructions for performing common tasks.	<ul style="list-style-type: none"> ■ The Documentation Library, which is available in the Symantec Management Console on the Help menu. ■ The Supported Products A-Z page, which is available at the following URL: http://www.symantec.com/business/support/index?page=products Open your product's support page, and then under Common Topics, click Documentation.

Table 1-3 Documentation resources (*continued*)

Document	Description	Location
Help	<p>Information about how to use this product, including detailed technical information and instructions for performing common tasks.</p> <p>Help is available at the solution level and at the suite level.</p> <p>This information is available in HTML help format.</p>	<p>The Documentation Library, which is available in the Symantec Management Console on the Help menu.</p> <p>Context-sensitive help is available for most screens in the Symantec Management Console.</p> <p>You can open context-sensitive help in the following ways:</p> <ul style="list-style-type: none"> ■ The F1 key when the page is active. ■ The Context command, which is available in the Symantec Management Console on the Help menu.

In addition to the product documentation, you can use the following resources to learn about Symantec products.

Table 1-4 Symantec product information resources

Resource	Description	Location
SymWISE Support Knowledgebase	Articles, incidents, and issues about Symantec products.	http://www.symantec.com/business/theme.jsp?themeid=support-knowledgebase
Symantec Connect	An online resource that contains forums, articles, blogs, downloads, events, videos, groups, and ideas for users of Symantec products.	http://www.symantec.com/connect/endpoint-management

Configuring the Monitor Solution Server

This chapter includes the following topics:

- [Configuring the monitor server](#)
- [About monitor packs](#)
- [Importing monitor packs](#)
- [About database maintenance](#)
- [Maintaining collected performance data](#)
- [About heartbeat](#)
- [Setting up the monitor server's heartbeat settings](#)

Configuring the monitor server

The following describes the process for preparing the monitor server.

Table 2-1 Process for configuring the monitor server

Step	Action	Description
Step 1	Import a monitor pack.	Monitor packs include the necessary monitor policies, metrics, rules, and tasks for monitoring an operating system or application. Monitor packs also contain preconfigured monitor policies with preset thresholds and severities. You must import a monitor pack to monitor computers and devices. See “ Importing monitor packs ” on page 22.

Table 2-1 Process for configuring the monitor server (*continued*)

Step	Action	Description
Step 2	Set up database maintenance.	Monitor Solution collects data from monitor computers and stores it in the database. You can configure the database maintenance settings to define when data is summarized and purged. See “ Maintaining collected performance data ” on page 25.
Step 3	Configure heartbeat monitoring settings.	Monitor Solution collects heartbeat signals from Monitor Plug-ins. You can configure the server-side heartbeat settings to define how often Monitor Solution checks for heartbeats. Specify the number of failures that are allowed to occur before Monitor Solution sends an alert to the Event Console. See “ Setting up the monitor server’s heartbeat settings ” on page 27.

About monitor packs

Monitor packs are available for monitoring many aspects of your computer resources and network to ensure their availability. Monitor packs include the necessary monitor policies, metrics, rules, and tasks for monitoring an operating system or application. Monitor packs also contain preconfigured monitor policies with preset thresholds and severities. You can import a monitor pack by scheduling a monitor pack import on the **Import Monitor Pack** page.

Monitor packs must be imported following the installation of Monitor Solution. Therefore, you can choose what functionality you want installed on your monitor server and when you want it to be installed. However, you must import a monitor pack before you can use Monitor Solution to monitor devices.

See “[Importing monitor packs](#)” on page 22.

Importing monitor packs

Monitor packs include the necessary monitor policies, metrics, rules, and tasks for monitoring an operating system or application. Monitor packs also contain preconfigured monitor policies with preset thresholds and severities.

You must import monitor packs following the installation of Monitor Solution. Importing monitor packs lets you choose what functionality you want installed

on your monitoring server, and when you want it to be installed. Importing a monitor pack is accomplished by scheduling a monitor pack import in the **Import Monitor Pack** page.

See “[Configuring the monitor server](#)” on page 21.

See “[About monitor packs](#)” on page 22.

You can also import monitor packs from the **Monitoring and Alerting** section of the **First Time Setup** portal. The **First Time Setup** portal is available on the **Home** menu, under Notification Server Management.

For more information page, see the topic about performing the First Time Setup configuration in the *Symantec Management Platform User Guide*.

To create an import policy to import a monitor pack

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Import Monitor Pack**.
- 3 On the **Import Monitor Pack** page, select the monitor pack to import.
- 4 Click **Schedule**.
- 5 In the **Schedule Monitor Pack** dialog box, choose from the following options:

Run now	Starts the import immediately after the dialog is completed.
----------------	--

Run on a schedule	Lets you specify a time and date for when you want the import to run.
--------------------------	---

Yield to system resources	Instructs the policy to aggressively consume system resources during import.
----------------------------------	--

Override existing items	Replaces any items that are already stored in the database.
--------------------------------	---

- 6 Click **OK** to apply the configuration settings to the policy and enable it to run.

About database maintenance

You have control over how the data that you gather from monitoring is handled. You can create a custom configuration for the summarization and purging of Monitor Solution data. The data that you choose to store can be as granular or as broad as you require. For example you can specify to store highly detailed and

granular data (typically in 5-minute intervals) for the current day, week, or month. You can specify to then keep less-detailed data that is summarized for the previous days, weeks, or months. Ultimately, the data is purged from the database entirely when it reaches a final age. This control helps you to store the data that you require while also helping you control database growth. However, because data becomes summarized according to a schedule, it also becomes less detailed and therefore less reliable. This factor should be taken into consideration when you schedule the summarization and purging schedule.

Schedule database purging to occur during non-peak times of the day. Large amounts of data can cause the summarization and purging process to take an extended amount of time.

Numeric data passes through the following four stages according to the time-periods that you define:

- When data is first collected it is stored as fully detailed data for as long as you require.
- After the specified time period lapses, detailed data is rolled-up into hourly summaries.
- After the specified time period lapses, hourly summaries are rolled-up into daily summaries.
- After the specified time period lapses, daily summaries are purged from the database. Database summarization and purging occurs daily at a time you define.

Non-numeric data (such as string metric, process, and NT event data) is not summarized. It is instead stored in full detail for as long as you define and then purged when the specified time period lapses.

The metric polling intervals of a metric can affect the amount of data that is collected. Monitor Plug-in settings and Remote Monitoring Server settings can also affect the amount of data that is stored. For example, you can choose to log NT Event data only when an alert raises. Changing this setting decreases the overall amount of NT Event data that is collected and therefore also decreases the amount of data that is stored.

See “[Maintaining collected performance data](#)” on page 25.

See “[Monitor Plug-in configuration settings: Data Collection tab](#)” on page 39.

See “[Remote Monitoring Server Settings: Data Collection tab](#)” on page 59.

Maintaining collected performance data

You can control when data summarization and purging occurs. This ability lets you maintain the performance data that you collect.

See “[Configuring the monitor server](#)” on page 21.

See “[About database maintenance](#)” on page 23.

To maintain collected performance data

- 1 In the Symantec Management Console, on the Home menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Settings > Monitor Server Settings**.
- 3 On the **Monitor Server Settings** page, click the **Purge Maintenance** tab.

4 Configure the following options:

Detailed data	How long you want to store detailed numeric performance data before it is purged and summarized into hourly values.
Hourly summaries	How long you want to store hourly summarizations of numeric performance data before it is purged and summarized into daily values. The value that is used is the Detailed data value plus the Hourly summaries value.
Daily summaries	How long you want to store daily summarizations of numeric performance data before it is purged from the database. The value that is used is the Detailed data value plus the Hourly summaries value plus the Daily summaries value.
String metric data	How long you want to store string data before it is purged from the database.
Process data	How long you want to store Process data before it is purged from the database.
NT event data	How long you want to store NT event data before it is purged from the database.
Command timeout	How much time should be allowed to pass before a purging timeout is declared.
Perform daily purge at	What time of the day the database purging should occur.

5 Click **Save changes**.

About heartbeat

The Monitor Plug-in can send scheduled messages to the Notification Server computer. These messages are called heartbeats. Heartbeats are used to determine if a Monitor Plug-in is available and in communication with the Notification Server computer. In addition to monitoring the health of a Monitor Plug-in, heartbeats also give insight into the host system's uptime and availability. The Monitor Plug-in is a service on the monitored computer so Monitor Solution equates uptime to the time that the Monitor Plug-in is up. A computer is more likely to be up and available if a Monitor Plug-in on that computer is running and communicating. Likewise, if expected heartbeats are not received from a Monitor Plug-in, it could

indicate a possible problem. The problem can be with either the computer that hosts the non-responsive plug-in or with the network connection that is used.

You can use heartbeat data to display reports on the availability of Monitor Plug-ins.

Should a heartbeat fail to be received, an alert is generated in the Event Console. With Event Console you can create rules to automatically execute server-side tasks in response to a failed heartbeat alert. For example, you can have a rule configured to send an email to you when a heartbeat fails. You can even configure the rule to execute a run script task that you preconfigure to diagnose the source of the heartbeat failure.

See “[Setting up the monitor server's heartbeat settings](#)” on page 27.

See “[Remote Monitoring Server Settings: Data Collection tab](#)” on page 59.

Setting up the monitor server's heartbeat settings

You can define heartbeat settings for the monitor server. These settings control how often the monitor server checks for received heartbeats.

See “[Configuring the monitor server](#)” on page 21.

See “[About heartbeat](#)” on page 26.

You can also define settings for how often the Monitor Plug-in sends heartbeats and whether-or-not it records system uptime. You can define these settings in the Monitor Plug-in configuration settings page.

See “[Monitor Plug-in configuration settings: Data Collection tab](#)” on page 39.

To set monitor server's heartbeat settings

- 1 In the Symantec Management Console , on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Settings > Monitor Server Settings**.
- 3 Click the **Heartbeat** tab.

4 Specify values for the following fields:

Check for heartbeat every	How often the monitor server should check for heartbeats.
Retry every	How often the server should attempt to confirm a heartbeat if an expected heartbeat is not received.
Retry attempts	How many attempts should be made to confirm that a heartbeat has been sent before the server generates an alert.

5 Click **Save changes**.

Configuring the Monitor Plug-in

This chapter includes the following topics:

- [About the Monitor Plug-in](#)
- [About Monitor Plug-in installation policies](#)
- [About Monitor Plug-in configuration policies](#)
- [About Monitor policies](#)
- [About Monitor Plug-in profiling](#)
- [Preparing managed computers for agent-based monitoring](#)
- [Creating new Monitor Plug-in settings](#)
- [Configuring Monitor Plug-in settings](#)
- [Monitor Plug-in configuration settings: General tab](#)
- [Monitor Plug-in configuration settings: Performance Tuning tab](#)
- [Monitor Plug-in configuration settings: Data Collection tab](#)
- [Monitor Plug-in configuration settings: Maintenance Windows tab](#)
- [Installing the Monitor Plug-in](#)
- [Upgrading the Monitor Plug-in](#)
- [Uninstalling the Monitor Plug-in](#)

About the Monitor Plug-in

The Monitor Plug-in is a monitoring application that you install on target computers so that you can monitor them. The Monitor Plug-in requires and works with the Symantec Management Agent to communicate with the Notification Server computer.

See “[Preparing managed computers for agent-based monitoring](#)” on page 33.

The Monitor Plug-in collects the following types of data:

- **Inventory**

This data is collected and sent to the Notification Server computer on a schedule. This data is used in reports and contains information about application detection and plug-in configuration on the monitored computer.

- **Performance**

This data is sent to the Notification Server computer on a different schedule than inventory data. Performance data is used in historical performance graphs by the Historical Performance Viewer and in reports. Some of this data can also be sent directly to Real-time Performance Viewer for real-time graphs.

See “[Viewing real-time performance data](#)” on page 133.

- **Alerts**

Rules can trigger alerts. The Monitor Plug-in generates alert data whenever a rule evaluation results in a change of state. Each rule can have a Normal, Informational, Undetermined, Warning, Major, or Critical state. The individual rule states are aggregated into an overall state for the computer. A resource can have an aggregated state of Normal, Warning, Major, or Critical.

To operate, Monitor Plug-in use the following three major types of policies:

- **Monitor Plug-in installation policies**

See “[About Monitor Plug-in installation policies](#)” on page 30.

- **Monitor Plug-in configuration policies**

See “[About Monitor Plug-in configuration policies](#)” on page 31.

- **Monitor policies**

See “[About Monitor policies](#)” on page 31.

About Monitor Plug-in installation policies

The Monitor Plug-in is installed on targeted computers through the use of a Monitor Plug-in installation policy. This installation policy installs the plug-in onto collections of computers according to a schedule that you define. Similar policies for upgrading and uninstalling the Monitor Plug-in are also available.

To install the Monitor Plug-in, you configure a policy that installs it on target computers. You select the group of computers on which the policy runs and schedule when it runs. If you choose a group that contains a computer that already has the plug-in installed, the task is ignored on that computer. When the policy is on, any computers that are added to the network, and are members of the group, automatically have the Monitor Plug-in installed.

See “[Installing the Monitor Plug-in](#)” on page 42.

Note: The Monitor Plug-in requires the Symantec Management Agent to perform tasks on the target computers and to communicate with the Notification Server computer. See “[Preparing managed computers for agent-based monitoring](#)” on page 33.

About Monitor Plug-in configuration policies

Monitor Plug-ins receive their configuration instructions from configuration policies. Configuration policies let you customize the settings of various Monitor Plug-ins for various purposes. The Monitor Plug-in can locally log performance data, process data, and NT event data.

You can configure if and how often these logs are stored locally and if and how often the data is uploaded to the CMDB. You can also specify how often local data is deleted. You can performance tune a Monitor Plug-in to meet the needs of your environment and define how the Monitor Plug-in performs during maintenance windows.

See “[Configuring Monitor Plug-in settings](#)” on page 35.

About Monitor policies

Monitor policies inform the Monitor Plug-in of what data you want monitored and how that data should be analyzed. The data is evaluated against the conditions of rules. Based on these rules the Monitor Plug-in can run automated actions in response to data that reaches an undesired state or range. The Monitor Plug-in returns the monitored data to the Notification Server computer. The Notification Server computer uses monitored data to run Task Server tasks, for real-time performance monitoring, and historical performance reporting.

Monitor policies specify application detection behavior for the Monitor Plug-in. Application detection enables the Monitor Plug-in to regularly check the monitored computer for the presence of applications that it has been configured to monitor. If an application that is able to be monitored is detected on the computer, the

plug-in automatically begins monitoring the application. If an application can be monitored but it is not installed or is removed, the plug-in does not attempt to monitor it. This functionality enables the Monitor Plug-in to dynamically adjust to the changing roles of a computer.

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

About Monitor Plug-in profiling

(Windows only)

Plug-in profiling is a rule triggering mechanism that triggers a rule-based on statistical criteria. Monitor Solution compares a current metric value with the average of previous values for the same metric. Monitor Solution determines if the current value is within a specified number of standard deviations from the average value for the metric. If the value is outside the range, the rule is triggered.

Note: Plug-in profiling is applicable only for metric-type rules.

Plug-in profiling lets Monitor Solution automatically and dynamically determine which metric values are acceptable and trigger a rule if necessary. As metric values change over time, the triggering range is automatically adjusted. Automatically adjusting the range helps prevent rules from being triggered when they should not be and reduces the need to manually adjust rules. This automatic adjustment also minimizes the need for manual adjustments across different computer hardware and software.

Plug-in profiling can be used along with any other rule triggering mechanisms associated with a rule. Plug-in profiling requires that the local logging of data be enabled (the local logging is automatically enabled when the plug-in profiling feature is enabled).

Warning: The plug-in profiling feature can be resource-intensive if polling intervals are small. It can also be resource-intensive if many metrics are involved and if a large number of days are used in the profiling.

A minimum of seven days of data is required to use plug-in profiling. More data can be used if wanted. The amount of data that is used is the amount of data logged, which is specified on the Plug-in Configuration page.

See “[Monitor Plug-in configuration settings: Data Collection tab](#)” on page 39.

Monitor Solution uses all of the locally logged data that is available for the current time block to calculate the average and the standard deviation. Each day is divided

into three-hour time blocks. The first time block begins at midnight. Monitor Solution uses all of the data that was collected between 12:00 A.M. and 3:00 A.M. for all Tuesdays if the current time is 12:30 A.M. on Tuesday. Monitor Solution then calculates the average and calculates the standard deviation to determine if the rule should be triggered.

For a normal distribution of data, which the data might or might not have, one standard deviation includes 68% of the data points. Two standard deviations include 95% of the points. Three standard deviations include 99.7% of data points. If you set the plug-in profiling to three standard deviations, the rule triggers 0.3% of the time. The normal distribution values are meant only as a rough guide.

Preparing managed computers for agent-based monitoring

Some monitor tasks require that the target computers be managed. Managed computers are the computers that have the Symantec Management Agent installed on them. The Monitor Plug-in is dependent upon the Symantec Management Agent. The Remote Monitoring Server provides limited monitoring functionality without a plug-in. Detailed monitoring requires the installation of both the Symantec Management Agent and the Monitor Plug-in.

For more information, see the topics about the Symantec Management Agent in the *Symantec Management Platform User Guide*.

To prepare managed computers for monitoring, you must complete the following steps.

Note: You had the opportunity to perform these steps at the time of installation or during the initial setup.

Table 3-1 Process for preparing managed computers for agent-based monitoring

Step	Action	Description
Step 1	Discover the computers that you want to manage.	When computers are discovered, resource objects are created for them in the CMDB. You may have discovered computers when you installed Notification Server or when you added new computers to the network. For more information, see the <i>Symantec Management Platform User Guide</i> .

Table 3-1 Process for preparing managed computers for agent-based monitoring (*continued*)

Step	Action	Description
Step 2	Manage the computers by installing the Symantec Management Agent.	You may have performed this step when you installed Notification Server or when you added new computers to the network. For more information, see the <i>Symantec Management Platform User Guide</i> . The Symantec Management Agent has two versions: one for Windows and one for UNIX, Linux, and MAC.
Step 3	Prepare managed computers by installing the Monitor Plug-in.	To monitor computers with agent-based monitoring, you must install the Monitor Plug-in on target computers. See “ Installing the Monitor Plug-in ” on page 42.

Creating new Monitor Plug-in settings

You can create new Monitor Plug-in settings within a Monitor Plug-in configuration policy. When a configuration policy is turned on, the settings apply to all the Monitor Plug-ins that are contained in the groups that the policy targets.

Note: If more than one configuration policy targets a Monitor Plug-in, the plug-in’s settings match the configuration policy that was saved most recently.

See “[About the Monitor Plug-in](#)” on page 30.

See “[Configuring Monitor Plug-in settings](#)” on page 35.

To create new Monitor Plug-in settings

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Agents/Plug-ins**.
- 3 Expand the folder for the operating system or application that you want the Monitor Plug-in to run on.

For example: **Windows**.

- 4 Right-click the **Configuration** folder and click **New > New Plug-in Settings**.
A new Monitor Plug-in configuration item is created under the **Configuration** folder with the name **New Plug-in Settings**.
- 5 To rename the plug-in configuration item, right-click **New Plug-in Settings**, and click **Rename**.
- 6 Rename the plug-in configuration item, and click **Save changes**.

Configuring Monitor Plug-in settings

The Monitor Plug-in settings are configured and applied to Monitor Plug-ins with configuration policies. When a configuration policy is turned on, the settings apply to all the Monitor Plug-ins that are contained in the groups that the policy targets.

See “[About the Monitor Plug-in](#)” on page 30.

See “[Creating new Monitor Plug-in settings](#)” on page 34.

See “[Installing the Monitor Plug-in](#)” on page 42.

Note: If more than one configuration policy targets a Monitor Plug-in, the plug-in’s settings match the configuration policy that was saved the most recently.

Within a configuration policy, you can configure the following Monitor Plug-in settings:

- Application detection
- Metric preferences
- Data collection settings
- Monitor Plug-in Heartbeat settings
- SNMP settings
- SQL settings
- Real-time Performance Viewer parameters
- Maintenance window settings

To configure Monitor Plug-in settings

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Agents/Plug-ins**.

- 3** Expand the folder for the operating system or application that the Monitor Plug-in runs on, and then expand the **Configuration** folder.

For example: **Windows > Configuration**.

- 4** Select the plug-in configuration item.

- 5** In the right pane, under **Plug-in Config Settings**, click the tabs to configure the following settings:

General See “[Monitor Plug-in configuration settings: General tab](#)” on page 36.

Performance Tuning See “[Monitor Plug-in configuration settings: Performance Tuning tab](#)” on page 38.

Data Collection See “[Monitor Plug-in configuration settings: Data Collection tab](#)” on page 39.

Maintenance Windows See “[Monitor Plug-in configuration settings: Maintenance Windows tab](#)” on page 42.

- 6** Under **Applies To**, apply the policy to groups of computers.

- 7** Click **Save changes**.

Monitor Plug-in configuration settings: General tab

This tab lets you configure general Monitor Plug-in settings within a Monitor Plug-in configuration policy.

See “[Configuring Monitor Plug-in settings](#)” on page 35.

Table 3-2 Settings on the **General** tab

Setting	Description
Monitor Policy Detection	<p>Run detection every – The frequency in minutes that the Monitor Plug-in checks for a new application, or removed application that it can monitor. The plug-in automatically enables or disables any rules that are associated with a detected or removed application. This ability ensures that newly installed applications are monitored in a timely manner and attempts are not made to monitor removed applications.</p> <p>See “About application detection” on page 68.</p>

Table 3-2 Settings on the **General** tab (*continued*)

Setting	Description
Performance Viewer	Maximum concurrent connections – The maximum number of connections from the Real-time Performance Viewer that are allowed at one time on a monitored computer. TCP/IP port – The port number that the Performance Viewer uses to connect and receive real-time metric values. The default port on the client computer is 1011. Use socket server authentication - If you receive an unauthorized request error message during an attempt to retrieve metric data from a Monitor Plug-in, uncheck this option. This option secures a connection between the Monitor Plug-in web service and the Monitor Plug-in by passing a token. This token passing function can fail if an environment has certain security implementations. When you uncheck this option, it disables this token passing function.
SNMP	Community string – The relationship between an SNMP server system and the client computer systems. This string acts as a password to control the client access to the server. Many devices use “public” as the default read-only community string and “private” as the default read-write community string.
SQL Metric Default Connection	Server name – The name of computer where SQL database is running. Database name – The name of the database. Connect using – The default authentication method that the Monitor Plug-in uses to access the database. The SQL Query Builder uses these settings to authenticate to the SQL database. The options are as follows: <ul style="list-style-type: none">■ Windows Authentication – If you select Windows authentication, you must specify a Windows account logon name and password to use to authenticate■ SQL Authentication – If you select SQL Server authentication, you must specify the SQL Server logon name and password.■ NS Authentication – If you select NS Authentication, the Notification Server computer authentication settings are used.

Monitor Plug-in configuration settings: Performance Tuning tab

This tab lets you configure general performance settings within a Monitor Plug-in configuration policy.

See “[Configuring Monitor Plug-in settings](#)” on page 35.

Table 3-3 Settings on the Performance Tuning tab

Setting	Description
Metrics	<p>Polling threads – The number of threads that the Monitor Plug-in metric provider component uses to read metric data. A scheduler queues polling requests for each metric. As a thread becomes available, the next metric in the queue is polled. If too few threads are allocated, the Monitor Plug-in does not poll as frequently as defined in the metric definitions. If too many threads are allocated and in use, there is an increase in resource usage on the Monitor Plug-in.</p> <p>Initialization interval – When the Monitor Plug-in first starts, the first polling time for each configured metric is scheduled. This value indicates the time in milliseconds between these initial polls (as opposed to attempting to do all initial polls at the same time). Staggering the initial polling prevents over-utilization of resources on startup.</p> <p>Unavailable metric notification – The notification methods that are used when a metric is unavailable. This error is recorded in the Monitor Solution audit log (in Evt_Monitor_Metric_Status table in Resource Manager). This error is also recorded in the CE_LOG file on the Notification Server computer.</p> <p>The options are as follows:</p> <ul style="list-style-type: none">■ Notification Server event – Generates an event that is sent to the Notification Server computer.■ NTEvent – Generates a Windows NT event that is added to the event log on a Windows server.

Table 3-3 Settings on the **Performance Tuning** tab (*continued*)

Setting	Description
Alerts	Alert batching – Specifies the amount of time alerts are batched before they are sent. Batching alerts lets the Monitor Plug-in send alerts more efficiently. If there are a large number of alerts in a very short time frame, this feature speeds up the alert sending process. When alerts are infrequent, set this option to a smaller value. Setting this option to a smaller value reduces the amount of time an alert needs to wait before it is sent.

Monitor Plug-in configuration settings: Data Collection tab

This tab lets you configure data collection and log file management settings within a Monitor Plug-in configuration policy.

See “[Configuring Monitor Plug-in settings](#)” on page 35.

Table 3-4 Settings on the **Data Collection** tab

Setting	Description
Data Collection	<p>Record metric values every – Specifies how often all the monitored metric values are time-stamped and recorded in the performance log. The performance log records only the values that have changed. Setting this value too low results in large logs. Empty time-stamp entries are written if this setting is smaller than the interval of the metric that is polled most frequently.</p> <p>Record process values every – Specifies how often process data is timestamped and recorded in the performance log. Increasing this value reduces the frequency of logging process data. In addition to recording the values at this interval, the performance log can be automatically updated when a rule is triggered.</p> <p>Record process values when alerts get triggered – Instructs the Monitor Plug-in to record process values when alerts are raised.</p> <p>Record NT event data – Controls if and when NT Event data is recorded. If this option is selected, the options are as follows:</p> <ul style="list-style-type: none"> ■ Always when referenced by a rule ■ Only when an alert is raised

Table 3-4 Settings on the **Data Collection** tab (*continued*)

Setting	Description
System Uptime	Send heartbeat every – Controls how often the Monitor Plug-in sends heartbeat signals to the Notification Server computer. See “ About heartbeat ” on page 26.

Table 3-4 Settings on the Data Collection tab (*continued*)

Setting	Description
Log File Management	<p>The Monitor Plug-in can store metric data in log files that are uploaded to the Notification Server for use in the Historical Performance viewer. Metric logs can also be compressed and stored locally for use with Monitor Plug-in profiling. When each metric log is started, the initial metric values are stored. Subsequent logs store only the deviations from the initially recorded metric values. For example, if a metric value does not change over a time period, then only the initially recorded value is stored.</p> <p>A log must be closed before it can be used. Once a metric log is completed, it is considered closed. Logs are closed automatically when the specified time completes, or when the log file size is too large.</p> <p>Close logs every – Controls how often metric log files are closed and new metric log files are started. Closed log files can be uploaded to Notification Server and copied locally for Monitor Plug-in profiling. Log files are automatically closed when the file size reaches 15 Mb.</p> <p>Increasing this value allows for better compression and reduces the overall amount of data that is stored in the database. However, increasing this value can also increase the amount of time that you must wait before the metric log is available.</p> <p>Upload logs to the Notification Server every – Controls how often closed metric logs are uploaded to the Notification Server.</p> <p>For example, if Close logs every is set to 60 minutes and Upload logs to the Notification Server every is set to 15 minutes, then the log is uploaded within 15 minutes after the log closes.</p> <p>Note: This setting does not affect the closing of metric logs. If a log is not closed, then it is not uploaded.</p> <p>Profile metric data in blocks of – Lets you specify a time period for compressing log files for use with Monitor Plug-in profiling.</p> <p>This option compresses data from closed metric logs into metric blocks. A metric block is the length of time for which values are averaged together into a single value for profiling. In order for plug-in profiling to work this option must be enabled. Plug-in profiling must also be enabled within any rules against which you want to perform plug-in profiling.</p> <p>See “About Monitor Plug-in profiling” on page 32.</p> <p>Save logs locally on the Monitor Plug-in for – Lets you save a local, compressed copy of the log data that was uploaded to Notification Server. This option must be enabled to use the plug-in profiling feature.</p> <p>Increasing this value increases the accuracy of profiling, but also increases the Monitor Plug-in’s footprint on the client computer.</p>

Monitor Plug-in configuration settings: Maintenance Windows tab

This tab lets you configure settings for maintenance windows within a Monitor Plug-in configuration policy.

Note: **Maintenance Windows** tab is only applicable to the Agent-based monitored resources.

See “[Configuring Monitor Plug-in settings](#)” on page 35.

Table 3-5 Settings on the **Maintenance Windows** tab

Setting	Description
Maintenance Windows	<p>Continue evaluating monitor policy rules and monitoring heartbeats - Instructs the Monitor Plug-in to continue monitoring during maintenance windows.</p> <p>If this setting is selected the options are as follows:</p> <ul style="list-style-type: none">■ If an alert is raised, change severity to informational - If selected, all alerts that are raised during maintenance windows have a severity setting of informational. See “About severity states” on page 119.■ Run tasks - If selected, Monitor Plug-ins continue to run actions and tasks, even during maintenance windows. See “About Monitor Solution tasks and actions” on page 117.

Installing the Monitor Plug-in

To install the Monitor Plug-in, you configure a policy that installs it on target computers. You select the group of computers on which the policy runs and schedule when it runs. If you choose a group that contains a computer that already has the Monitor Plug-in installed, the task is ignored on that computer.

See “[About the Monitor Plug-in](#)” on page 30.

See “[Upgrading the Monitor Plug-in](#)” on page 43.

See “[Uninstalling the Monitor Plug-in](#)” on page 44.

Before performing this task, you must install the Symantec Management Agent on target computers.

See “[Preparing managed computers for agent-based monitoring](#)” on page 33.

Note: Monitor Solution has separate Plug-in rollout policies for 32-bit computers and 64-bit computers.

You can also install the Monitor Plug-in for each of the supported operating systems from the **Monitoring and Alerting** section of the **First Time Setup** portal. The **First Time Setup** portal is available on the **Home** menu, under Notification Server Management.

For more information page, see the topic about performing the First Time Setup configuration in the *Symantec Management Platform User Guide*.

To install the Monitor Plug-in

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Agents/Plug-ins**.
- 3 Expand the folder for the operating system or application that the Plug-in is designed to run on. Expand the **Rollout** folder, and select the applicable installation policy.
For example: click **Windows > Rollout > Monitor Plug-in for Windows x86 - Install**.
- 4 On the **Monitor Plug-in for Windows x86 - Install** page, turn on the policy. At the upper right of the page, click the colored circle and click **On**.
- 5 Click **Apply to** to select the computers to install the plug-in on.
In most cases, you can use the default group to install the Plug-in on all eligible computers that do not have it installed.
- 6 Schedule the policy.
- 7 Click **Save changes**.

Upgrading the Monitor Plug-in

To upgrade the Monitor Plug-in, you configure a policy that installs and upgrades it on target computers. You select the group of computers on which the policy runs and schedule when it runs. If the group contains a computer that already has the correct version of the Monitor Plug-in installed, the task is ignored on that computer.

See “[About the Monitor Plug-in](#)” on page 30.

See “[Installing the Monitor Plug-in](#)” on page 42.

See “[Uninstalling the Monitor Plug-in](#)” on page 44.

Note: This upgrade policy does not upgrade version 6.x of the Monitor Plug-in for UNIX and Linux client computers. For these computers, use the install policy to deploy the Monitor Plug-in. You can use this upgrade policy to upgrade version 7.0 or 7.1 of the Monitor Plug-in for UNIX and Linux Client computers to version 7.1 SP2.

To upgrade the Monitor Plug-in

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Agents/Plug-ins**.
- 3 Expand the folder for the operating system or application that the plug-in is designed to run on. Expand the **Rollout** folder, and select the applicable upgrade policy.
For example: expand **Windows > Rollout > Monitor Plug-in for Windows x86 - Upgrade**.
- 4 On the **Monitor Plug-in for Windows x86 - Upgrade** page, turn on the policy. At the upper right of the page, click the colored circle and click **On**.
- 5 On the **Monitor Plug-in for Windows x86 - Upgrade** page, click **Apply to** to select the computers to upgrade the plug-in on.
- 6 Schedule the policy.
- 7 Click **Save changes**.

Uninstalling the Monitor Plug-in

You use a policy to uninstall the Monitor Plug-in. You select the group of computers on which the policy runs and schedule when it runs. If you choose a group that contains a computer that does not have the Monitor Plug-in installed, the task is ignored on that computer.

See “[About the Monitor Plug-in](#)” on page 30.

See “[Installing the Monitor Plug-in](#)” on page 42.

See “[Upgrading the Monitor Plug-in](#)” on page 43.

To uninstall the Monitor Plug-in

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Agents/Plug-ins**.
- 3 Expand the folder for the operating system or application that the plug-in is designed to run on. Expand the **Rollout** folder and select the applicable uninstall policy.
For example: **Windows > Rollout > Monitor Plug-in for Windows x86 - Uninstall**.
- 4 On the **Monitor Plug-in for Windows x86 - Uninstall** page, click **Apply to** to select the computers to uninstall the plug-in from.

Warning: By default the uninstall policy is targeted to **All computers with Monitor Plug-in Installed**. If the **Apply to** option is not changed, turning on the policy uninstalls the Monitor Plug-in from all computers.

- 5 On the **Monitor Plug-in for Windows x86 - Uninstall** page, turn on the policy. At the upper right of the page, click the colored circle and click **On**.
- 6 Schedule the policy.
- 7 Click **Save changes**.

Configuring agentless monitoring

This chapter includes the following topics:

- [About agentless monitoring](#)
- [About agentless monitoring and network discovery](#)
- [About monitor service](#)
- [Setting up a remote monitoring site server](#)
- [Configuring remote monitoring server settings](#)

About agentless monitoring

You use agentless monitoring to monitor the computers that do not have the Monitor Plug-in. You monitor these computers with agentless monitoring policies. Because the Monitor Plug-in is not available on the computer, fewer aspects of the computer are available to be monitored. You use monitor service on a site server to perform agentless monitoring.

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

See “[About monitor service](#)” on page 48.

All agentless monitoring policies have a list of resource targets. These resource targets are the resources that are monitored. Each monitor service monitors the resources that its site server is assigned if an agentless monitoring policy targets those resources. Consequently, multiple site servers can monitor the same resource that is targeted in an agentless monitor policy. Also, different site servers can monitor the different resources that are targeted in the same agentless monitor policy.

You use agentless monitoring for the following reasons:

- You cannot install the Symantec Management Agent on the device that you want to monitor.
For example, VMware recommends that you not run third-party software in the VMware ESX Server service console. Another example would be a device that has an embedded system.
- You want to monitor the availability of a server.
In most cases you need to use agentless monitoring to perform an availability (ping) monitor.

About agentless monitoring and network discovery

A dependency exists between network discovery and agentless monitoring. The connection profiles within the network discovery bind those computers that are specified within the selected IP range. This association is required for the resources that use both agentless monitoring and any of the following metric sources:

- HTTP
- SNMP
- WMI
- WS-MAN

However if you want to monitor the availability status of a resource that use for ICMP, you do not need to run the network discovery task.

See “[About agentless monitoring](#)” on page 47.

For more information about connection profiles, see the *Symantec Management Platform User Guide*.

For more information about discovering network devices, see the *Symantec Management Platform User Guide*.

About monitor service

Monitor service on a site server lets you perform agentless monitoring. Monitor service is installed on the Notification Server computer by default.

See “[About agentless monitoring](#)” on page 47.

Because monitoring can be resource-intensive, you can distribute the monitoring load to other site servers to reduce the load on Notification Server. You can also

remove monitor service from the Notification Server computer to further reduce the load on this server.

See “[Setting up a remote monitoring site server](#)” on page 49.

Monitor service is integrated with the site server infrastructure. This integration lets the user specify the resources that each site server monitors.

Setting up a remote monitoring site server

You use monitor service on a site server to perform agentless monitoring. By default, monitor service is installed on the Notification Server computer. Because monitoring can be resource-intensive, you can distribute the monitoring load to other site servers to reduce the load on Notification Server. You can set up as many monitoring site servers as needed.

See “[About agentless monitoring](#)” on page 47.

See “[About monitor service](#)” on page 48.

To install the monitor service on a remote site server, the server must be running one of the following operating systems:

- Microsoft Windows Server 2003 SP2 x86
- Microsoft Windows Server 2008 R2 x64

The **Potential Monitor Servers** filter automatically determines possible site servers. The **Potential Monitor Servers** filter is available on the **Management** menu, under Filters. Your agentless monitor policies do not require any special configuration to work with a monitor service on one or more site servers.

Warning: You should only install monitor service on a computer that is secure and trusted. The security that is set up for the Notification Server computer must also apply to the site server computer.

Monitor service requires that the following be installed on the site server:

- The Symantec Management Agent
- The Pluggable Protocols Architecture (PPA) client computer component
- The credential manager client computer component

Table 4-1 Process for setting up a remote monitoring site server

Step	Action	Description
Step 1	Install the Symantec Management Agent on the site server.	A remote monitoring server and its dependencies require the Symantec Management Agent. If the Symantec Management Agent is not installed on the site server, install it. For more information, see topics about the Symantec Management Agent in the <i>Symantec Management Platform User Guide</i> .
Step 2	Configure connection profiles on Notification Server.	Connection profiles must be configured on Notification Server for remote monitoring to work. Configure your connection profiles before you install the Pluggable Protocols Architecture (PPA) client computer component on the site server. For more information, see topics about connection profiles in the <i>Symantec Management Platform User Guide</i> .
Step 3	Install the Pluggable Protocols Architecture (PPA) client computer component on the site server.	A remote monitoring server depends on the Pluggable Protocols Architecture (PPA) client computer component to communicate with network devices and computers. When the Pluggable Protocols Architecture (PPA) client computer component is installed, the credential manager client computer component is also installed. See “ Installing the Pluggable Protocols Architecture (PPA) client computer component on a site server ” on page 51.
Step 4	Add monitor service to one or more site servers.	You use the Site Management page to add monitor service to a site server. See “ Adding monitor service to a site server ” on page 53.

Table 4-1 Process for setting up a remote monitoring site server (*continued*)

Step	Action	Description
Step 5	(Optional) Remove monitor service from Notification Server.	You can remove monitor service from Notification Server to reduce the load on this server. See “ Removing monitor service from a site server ” on page 52.
Step 6	Configure the remote monitoring server settings.	You can configure the remote monitoring server settings. These are the global settings that apply to all monitor site servers. See “ Configuring remote monitoring server settings ” on page 55.
Step 7	(Optional) View the monitor site server reports.	The monitor site server reports let you determine which site servers monitor the resources that your agentless monitor policies target. See “ Monitor site server reports ” on page 54.

Installing the Pluggable Protocols Architecture (PPA) client computer component on a site server

Pluggable Protocols Architecture (PPA) includes a policy that can remotely install the Pluggable Protocols Architecture (PPA) client computer component on a site server. You must install this component on a site server before you can add monitor service to the site server. When the Pluggable Protocols Architecture (PPA) client computer component is installed, the credential manager client computer component is also installed. The policy that installs the credential manager client computer component configures the agent to automatically import credentials from Notification Server.

See “[Setting up a remote monitoring site server](#)” on page 49.

Warning: You should only install monitor service on a computer that is secure and trusted. The security that is set up for the Notification Server computer must also apply to the site server computer.

To install the Pluggable Protocols Architecture (PPA) client computer component on a site server

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, under **Settings**, click **Monitoring and Alerting > Protocol Management**.
- 3 Expand the **Protocol Management** folder and click **Install x86 Pluggable Protocols Agent Package** or **Install x64 Pluggable Protocols Agent Package**.
- 4 On the **Install Pluggable Protocols Agent Package** page, complete the following:
 - In the **Applied to** section, apply the policy to the site server.
For more information, see topics about specifying the targets of a policy in the *Symantec Management Platform User Guide*.
 - In the **Schedule** section, schedule when and how you want the policy to run.
For more information, see topics about specifying a policy schedule in the *Symantec Management Platform User Guide*.
 - Turn on the policy.
At the upper right of the page, click the colored circle and then click **On**.

5 Click **Save changes**.

After Pluggable Protocols Architecture (PPA) and credential manager are installed, wait until the Symantec Management Agent sends inventory information before adding a monitor service. You can confirm that the inventory information was sent on the **Symantec Management Agent Settings** tab of the Symantec Management Agent user interface on the site server.

Removing monitor service from a site server

You use monitor service on a site server to perform agentless monitoring. Monitor service is installed on the Notification Server computer by default. To reduce the load on Notification Server, you can remove monitor service from this server. You can also remove monitor service from any other site server.

See “[Setting up a remote monitoring site server](#)” on page 49.

To remove monitor service from a site server

- 1 In the Symantec Management Console, on the **Settings** menu, click **Notification Server > Site Server Settings**.
- 2 In the right pane, in the **Detailed Information** section, in the **View** menu, click **Site Servers**.
- 3 Select the site sever and click the **Edit** symbol.
- 4 In the **Add/Remove Services** dialog box, uncheck **Monitor Service** and click **Next**.
- 5 In the dialog box that appears, confirm your removal of monitor service from the correct site server and click **OK**.

Adding monitor service to a site server

You use monitor service on a site server to perform agentless monitoring. Monitor service is installed on the Notification Server computer by default. You can also add monitor service to one or more site servers.

See “[About agentless monitoring](#)” on page 47.

See “[About monitor service](#)” on page 48.

Before you can add monitor service to a site server, the following components must be installed on that server:

- Symantec Management Agent
- Pluggable Protocols Architecture (PPA) client computer component
- Credential manager client computer component
 - Credential manager client computer component is installed when you install Pluggable Protocols Architecture (PPA) client computer component. After Pluggable Protocols Architecture (PPA) and credential manager are installed, wait until the Symantec Management Agent sends inventory information before adding a monitor service.

See “[Setting up a remote monitoring site server](#)” on page 49.

Warning: You should only install monitor service on a computer that is secure and trusted. The security that is set up for the Notification Server computer must also apply to the site server computer.

When you add monitor service to a site server, it is installed on the selected site server according to the schedule in the installation policy. Monitor service has an installation policy for 64-bit and 32-bit computers. The installation policies

are in the **Advanced** folder for monitor service. To access these installation policies, on the **Settings** menu, click **Notification Server > Site Server Settings > Monitor Service > Advanced**.

To add monitor service to a site server

- 1 In the Symantec Management Console, on the **Settings** menu, click **Notification Server > Site Server Settings**.
- 2 In the right pane, in the **Detailed Information** section, in the **View** menu, click **Site Servers**.
- 3 Select the site server and click the **Edit** symbol.
- 4 In the **Add/Remove Services** dialog box, check **Monitor Service** and click **Next**.

If Pluggable Protocols Architecture (PPA) and credential manager are not installed on the site server, you cannot check **Monitor Service**.

- 5 In the dialog box that appears, confirm your addition of the monitor service to the correct site server and click **OK**.
- 6 To check the status of the installation, on the **Site Management** page, expand the Site Services section and then expand the Monitor Service section.

A pie chart displays the site servers that are installed, pending installation, or not installed.

Monitor site server reports

Monitor service on a site server lets you run agentless monitoring policies to monitor the resources that do not have the Symantec Management Agent installed. The monitor site server reports let you determine which site servers monitor the resources that your agentless monitor policies target.

See “[About agentless monitoring](#)” on page 47.

See “[About monitor service](#)” on page 48.

To access these reports, on the **Reports** menu, click **All Reports**, and then under **Reports**, click **Monitoring and Alerting > Monitor > Configuration > Monitor site server**.

The monitor site sever reports are as follows:

- **Monitored resources by RMS**

This report lists the resources that the selected site server monitors.

- **Resources not monitored by RMS**

This report lists the resources that no site server monitors.

■ **RMS by Monitored resources**

This report lists the site servers that monitor the selected resource.

Configuring remote monitoring server settings

You can configure the settings for the remote monitoring servers. You use a remote monitoring server and a monitor service to perform agentless monitoring.

See “[About agentless monitoring](#)” on page 47.

See “[About monitor service](#)” on page 48.

The remote monitoring server settings are the global settings that apply to all monitor site servers.

See “[Setting up a remote monitoring site server](#)” on page 49.

See “[Adding monitor service to a site server](#)” on page 53.

To configure remote monitoring server settings

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Settings > Remote Monitoring Server Settings**.
- 3 On the **Remote Monitoring Server Settings** page, click the tabs to configure the following settings:

General	See “ Remote Monitoring Server Settings: General tab ” on page 55.
----------------	--

Performance Tuning	See “ Remote Monitoring Server Settings: Performance Tuning tab ” on page 57.
---------------------------	---

Data Collection	See “ Remote Monitoring Server Settings: Data Collection tab ” on page 59.
------------------------	--

- 4 Click **Save changes**.

Remote Monitoring Server Settings: General tab

This tab lets you configure the general settings for the remote monitoring servers. You use a remote monitoring server and a monitor service to perform agentless monitoring. The remote monitoring server settings are the global settings that apply to all monitor site servers.

See “[Configuring remote monitoring server settings](#)” on page 55.

See “[About agentless monitoring](#)” on page 47.

Table 4-2 Settings on the **General** tab

Setting	Description
Monitor Policy Detection	<p>Run detection every – The frequency in minutes that the remote monitoring server checks for a new application or a removed application that it can monitor. The remote monitoring server automatically enables or disables any rules that are associated with a detected or removed application. This ability ensures that newly installed applications are monitored in a timely manner and attempts are not made to monitor removed applications.</p> <p>Note: This option is only applicable to detection on the remote monitoring servers. This option does not enable detection on the computers that are monitored through a remote monitoring server.</p> <p>See “About application detection” on page 68.</p>
Performance Viewer	<p>Maximum connections – The maximum number of connections from the Real-time Performance Viewer that are allowed at one time on a monitored computer.</p> <p>TCP/IP port – The port number that the Performance Viewer uses to connect and receive real-time metric values. The default port is 1011.</p> <p>Use socket server authentication – If you receive an unauthorized request error message during an attempt to retrieve metric data from a Monitor Plug-in, uncheck this option. This option secures a connection between the Monitor Plug-in web service and the Monitor Plug-in by passing a token. This token passing function can fail if an environment has certain security implementations. When you uncheck this option, it disables this token passing function.</p>
SNMP	<p>Community string – The relationship between an SNMP server system and the client computer systems. This string acts as a password to control the client access to the server. Many devices use “public” as the default read-only community string and “private” as the default read-write community string.</p> <p>Note: This option is only applicable in the case that an Agent-based policy is assigned to a remote monitoring server.</p>

Table 4-2 Settings on the **General** tab (*continued*)

Setting	Description
SQL Metric Default Connection	<p>Server name – The name of the computer where the SQL database is running.</p> <p>Database name – Name of the database.</p> <p>Connect using – The default authentication method that a remote monitoring server uses to access the database. The SQL Query Builder uses these settings to authenticate to the SQL database.</p> <p>The options are as follows:</p> <ul style="list-style-type: none">■ Windows Authentication – If you select Windows authentication, you must specify a Windows account logon user name and password to use to authenticate.■ SQL Authentication – If you select SQL Server authentication, you must specify the SQL Server logon user name and password.■ NS Authentication – If you select NS Authentication, the Notification Server computer authentication settings are used. <p>Note: This option is only applicable in the case that an Agent-based policy is assigned to a remote monitoring server.</p>

Remote Monitoring Server Settings: Performance Tuning tab

This tab lets you configure the performance settings for the remote monitoring servers. You use a remote monitoring server and a monitor service to perform agentless monitoring. The remote monitoring server settings are the global settings that apply to all monitor site servers.

See “[Configuring remote monitoring server settings](#)” on page 55.

See “[About agentless monitoring](#)” on page 47.

Table 4-3

Settings on the Performance Tuning tab

Setting	Description
Metrics	<p>Polling threads – The number of threads that the metric provider component of a remote monitoring server uses to read metric data. A scheduler queues polling requests for each metric. As a thread becomes available, the next metric in the queue is polled. If too few threads are allocated, the remote monitoring server does not poll as frequently as defined in the metric definitions. If too many threads are allocated and in use, there is an increase in resource usage in the remote monitoring server.</p> <p>Note: Increasing the number of status threads is useful in an environment with 1000 managed devices. For larger configurations, this may not change the performance.</p> <p>Initialization interval – When agentless monitoring first starts, the first polling time for each configured metric is scheduled. This value indicates the time in milliseconds between these initial polls (as opposed to attempting to do all initial polls at the same time). Staggering the initial polling prevents over-utilization of resources on startup.</p> <p>Unavailable metric notification – Check boxes for selecting the notification methods that are used when a metric is unavailable.</p> <p>The options are as follows:</p> <ul style="list-style-type: none"> ■ Notification Server event – Generates an event. ■ NTEvent – Generates a Windows NT event that is added to the event log on a Windows server.
Agentless metrics	<p>This section lets you determine how often credentials and connections details for agentless targets are refreshed.</p> <p>Refresh session every – Lets you control how often sessions are refreshed.</p> <p>Retry failed connections every – Lets you control how often attempts are made to re-establish a failed connection.</p> <p>Retry busy connections every – Lets you control how often attempts are made to establish a connection when the connection is busy.</p>

Table 4-3 Settings on the **Performance Tuning** tab (*continued*)

Setting	Description
Alerts	Alert batching – Specifies the amount of time alerts are batched before they are sent. If there are a large number of alerts in a very short time frame, this feature speeds up the alert sending process. When alerts are infrequent, set this option to a smaller value. Setting this option to a smaller value reduces the amount of time an alert needs to wait before it is sent.

Remote Monitoring Server Settings: Data Collection tab

This tab lets you configure how the remote monitoring server collects data. You use a remote monitoring server and a monitor service to perform agentless monitoring. The remote monitoring server settings are the global settings that apply to all monitor site servers.

See “[Configuring remote monitoring server settings](#)” on page 55.

See “[About agentless monitoring](#)” on page 47.

Table 4-4

Settings on the Data Collection tab

Setting	Description
Data Collection	<p>Record metric values every – Specifies how often all the monitored metric values are time-stamped and recorded in the performance log. The performance log records only the values that have changed. Setting this value too low results in large logs. Empty time-stamp entries are written if this setting is smaller than the interval of the metric that is polled most frequently.</p> <p>Record process values every – Specifies how often process data is time stamped and recorded in the performance log. Increasing this value reduces the frequency of logging process data. In addition to recording the values at this interval, the performance log is also automatically updated when a rule is triggered.</p> <p>Record process values when alerts get triggered – Instructs the remote monitoring servers to record process values when alerts are raised.</p> <p>Record NEvent data – Controls if and when NT Event data is recorded.</p> <p>If this option is selected, the options are as follows:</p> <ul style="list-style-type: none"> ■ Always when referenced by a rule ■ Only when an alert is raised
System Uptime	<p>Send heartbeat every – Controls how often the remote monitoring servers send heartbeat signals to the Notification Server computer.</p> <p>See “About heartbeat” on page 26.</p> <p>Record system uptime data – Instructs the remote monitoring servers to record system uptime.</p> <p>Note: These options are only applicable in the case that an Agent-based policy is assigned to a remote monitoring server.</p>

Table 4-4 Settings on the **Data Collection** tab (*continued*)

Setting	Description
Log File Management	<p>The log file management options are as follows:</p> <ul style="list-style-type: none">■ Close logs every – When a log is created, the initial values for all metrics being monitored are stored. Subsequently, only metric changes are stored. Increasing this value minimizes the amount of initial data that is stored. Increasing this value also allows for better compression because many of these strings are repeated. Increasing the value thereby reduces the overall amount of data that is stored in the database (even though individual logs can be larger in size). Log files are automatically closed when the file size reaches 1 MB.■ Upload logs to the Notification Server every – Check this option to upload performance logs to the CMDB. If data is not uploaded, you cannot view historical data in the Performance Viewer. If this option is selected, it specifies how often the performance log is closed and uploaded to the CMDB in minutes. The more frequently you upload data, the sooner the data is available for reports.■ Profile metric data in blocks of – Check this option to enable the metric data profiling feature. This feature enables the saving of logs locally. This feature must be enabled here and on the rule for which you want to perform plug-in profiling for the plug-in profiling feature to work. If this option is selected, it specifies the length of time for which values are averaged together into a single value for profiling. See “About Monitor Plug-in profiling” on page 32.■ Save logs locally on the Monitor plug-in for – Select this check box to save a copy of the uploaded data locally. A compressed local copy of the log is saved after the log is uploaded. This option must be enabled to use the plug-in profiling feature.

Working with Monitor Policies

This chapter includes the following topics:

- [Creating a monitor policy with the monitor policy wizard](#)
- [Creating a monitor policy](#)
- [Editing agent-based monitor policies](#)
- [Editing agentless monitor policies](#)
- [Adding rules to a monitor policy](#)
- [About application detection](#)
- [Adding application detection to a monitor policy](#)
- [Application detection types](#)
- [How operators work in application detection](#)
- [Adding computers to a monitor policy](#)

Creating a monitor policy with the monitor policy wizard

The monitoring of computers is accomplished by creating and applying a monitor policy. Monitor Solution includes a wizard that simplifies the process of creating monitoring policies. You can also create a monitor policy without the wizard.

See “[Creating a monitor policy](#)” on page 65.

Monitor policies use metrics, rules, and tasks to define the following information:

- What computer resources you want to have monitored?
- What metric data you want to be monitored?
- What fluctuations in the metric data imply about the status of the resource?
- What actions you want to occur when metric data reaches certain values?

You can also create a monitor policy with the monitor policy wizard from the **Monitoring and Alerting** section of the **First Time Setup** portal. The **First Time Setup** portal is available on the **Home** menu, under Notification Server Management.

For more information page, see the topic about performing the First Time Setup configuration in the *Symantec Management Platform User Guide*.

To create a monitor policy with the monitor policy wizard

- 1 In the Symantec Management Console, on the **Actions** menu, click **Monitor > New Policy**.
- 2 In the wizard, in the **Choose what to monitor** panel, choose what to monitor, and then click **Next**.

You must enter a name of the monitoring policy. A descriptive name can help you easily identify the policy in the future.

You must specify if the policy should be either an agent-based policy or an agentless policy. Whether or not the Monitor Plug-in is installed on the computers that you monitor determines if the policy should be agent-based or agentless.

- **Agent-based monitor** policies are intended to be run on computers with the Monitor Plug-in installed on them. If a computer has a Monitor Plug-in that is installed on it, more aspects of the computer are available to be monitored.
- **Agentless monitor** policies let you monitor computers without the Monitor Plug-in. Because the Monitor Plug-in is not available, fewer aspects of the computer are available to be monitored.

- 3 In the wizard, in the **Select monitoring categories** panel, select one or more monitoring categories, and then click **Next**.

A category is a grouping of rules. Rules are grouped into categories so that it is easier to organize and locate them. All of the rules that are contained in a category are displayed in the next step of the wizard.

- 4 In the wizard, in the **Add/Remove monitor rule** panel, add or remove monitor rules, and then click **Next**.

All of the rules of previously selected categories display in the window. Use **Add** and **Remove** to configure the policy to keep the rules that you want to have included in the monitoring policy.

See “[Adding rules to a monitor policy](#)” on page 68.

- 5 In the wizard, in the **Set rule actions** panel, set rule actions, and then click **Next**.

Based on the conditions of their rules, monitor policies are in one of six severity states. In this step of the wizard, you can specify the tasks that you want to occur for each severity state. Task server tasks are run from the task server. Monitor Plug-in tasks are run locally on the monitored computer. Tasks are run in the order they are displayed in the window.

See “[Adding actions to monitor policies](#)” on page 125.

- 6 In the wizard, in the **Select group of computer to monitor** panel, select a group of computers to monitor, and then click **Finish**.

Use **Apply to** to add the computer resources that you want the monitor policy to run on.

See “[Adding computers to a monitor policy](#)” on page 72.

- 7 Click **Finish**.

Creating a monitor policy

The monitoring of computers is accomplished by creating and applying a monitor policy. You can also create a monitor policy with the monitor policy wizard.

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

Monitor policies use metrics, rules, and tasks to define the following information:

- What computer resources you want to have monitored?
- What metric data you want to be monitored?
- What fluctuations in the metric data imply about the status of the resource?
- What actions you want to occur when metric data reaches certain values?

To create a monitor policy

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Monitor Policies**, and select the folder where you want to create the policy.
- 3 Right-click the folder, and then click **New > Monitor Policy (Agentless)** or **New > Monitor Policy (Agent-based)**.
- 4 Use the **Rules** tab to add, remove, and edit the rules that are contained in the monitor policy.
See “[Adding rules to a monitor policy](#)” on page 68.
- 5 (Agent-based policy only) Use the **Detection** tab to add, remove, and edit application detection settings for the monitoring policy.
See “[Adding application detection to a monitor policy](#)” on page 69.
- 6 Use the **Actions** tab to configure the tasks that you want the policy to run when different severities are reached.
See “[Adding actions to monitor policies](#)” on page 125.
- 7 Use the **Applies To** section to add and remove groups of the computers that you want the monitor policy to monitor.
See “[Adding computers to a monitor policy](#)” on page 72.
- 8 Click **Save changes**.

Editing agent-based monitor policies

You can activate, deactivate, and modify your agent-based monitoring policies.

To edit an agent-based monitoring policy

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Monitor Policies**, and select a policy.
- 3 On the policy page, turn on the policy.
On the upper right of the page, click the colored circle, and then click **On**.
- 4 Use the **Rules** tab to add, remove, and edit the rules that are contained in the monitor policy.
See “[Adding rules to a monitor policy](#)” on page 68.

- 5 Use the **Detection** tab to add, remove, and edit application detection settings for the monitoring policy.
See “[Adding application detection to a monitor policy](#)” on page 69.
- 6 Use the **Actions** tab to configure the tasks that you want the policy to run when different severities are reached.
See “[Adding actions to monitor policies](#)” on page 125.
- 7 Use the **Applies To** section to add and remove groups of the computers that you want the monitor policy to monitor.
See “[Adding computers to a monitor policy](#)” on page 72.
- 8 Click **Save changes**.

Editing agentless monitor policies

You can activate, deactivate, and modify your agentless monitoring policies.

To edit an agentless monitoring policy

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Monitor Policies**, and select a policy.
- 3 In the policy page, turn on the policy.
On the upper right of the page, click the colored circle, and then click **On**.
- 4 Use the **Rules** tab to add, remove, and select the rules that are contained in the monitor policy.
See “[Adding rules to a monitor policy](#)” on page 68.
- 5 Use the **Actions** tab to configure the tasks that you want the policy to run when the policy reaches different severities.
See “[Adding actions to monitor policies](#)” on page 125.
- 6 Use the **Monitored Targets** section to add and remove the computers and devices that you want the monitor policy to monitor.
See “[Adding computers to a monitor policy](#)” on page 72.
- 7 Click **Save changes**.

Adding rules to a monitor policy

You can add rules to a monitoring policy on the **Select Rule** page. You can also create, edit, and clone rules from this page.

See “[Creating and editing rules](#)” on page 110.

To add a rule to a monitor policy

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Monitor Policies**, and select a policy.
- 3 On the **Rules** tab, click the **Add** symbol.
- 4 Locate a rule and select it.
 - Use the **Search** field to search for a rule.
 - Use the **View** drop-down list to display rules within a specific category.
- 5 Select a rule, or use the CTRL key or SHIFT key to select multiple rules.
- 6 Click **OK**.

About application detection

Application detection searches a monitored computer to determine if specific applications are installed. When a Monitor policy runs the application detection methods, the computers that the policy applies to are searched to determine if specific applications are installed. If the applications are found, Monitor Solution knows to start monitoring. If the applications are not found, Monitor Solution knows to stop monitoring.

In a Monitor policy, you can specify which computers the policy applies to. Specifying which computers the policy applies to lets you save computer resources by targeting the most appropriate computers to run a particular policy on. In addition to filtering the monitored computers, a Monitor policy can also contain application detection methods.

When application detection methods are configured for a Monitor policy, the computers that the Monitor policy applies to are filtered at a more granular level. For example, you can configure a policy to monitor data when application X is installed. In that policy, you can specify that the policy applies to computers A through F. If application detection determines that application X is not on computer C, then the Monitor policy does not run on computer C. Monitor Solution does not attempt to collect monitoring data from computer C because the

application in question is not installed. Application detection saves system resources by eliminating the polling of metrics sources that are not available on the system.

When application detection is configured properly it helps you meet your monitoring goals in the following ways:

- Monitors the information and resources that are important to you.
- Uses the system resources more efficiently because you do not gather monitor data for any applications that are not installed on the computer.

When you install monitor packs , application detection is already configured for those packs. You can create additional application detection methods to add to an existing monitor policy. You can add additional application detection methods when you create customized categories or Monitor policies. Adding new application detection methods to your Monitor policies can help you meet any additional customization needs you may have.

See “[Adding application detection to a monitor policy](#)” on page 69.

You can choose from several application detection types. The different types let you monitor the specific applications that are the most critical to your operations.

See “[Application detection types](#)” on page 70.

Adding application detection to a monitor policy

Application detection searches a monitored computer to determine if specific applications are installed. When a monitor policy runs the application detection methods, the computers that the policy applies to are searched to determine if specific applications are installed. Monitor Solution can then start or stop monitoring those applications.

See “[About application detection](#)” on page 68.

To add detection to a monitor policy

- 1 Create or edit an agent-based monitor policy.

To create a monitor policy

See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

To edit an agent-based monitor policy

See “[Editing agent-based monitor policies](#)” on page 66.

- 2 On the **Detection** tab, click the **New** symbol to add a detection item.

- 3 Select an application detection type.

See “[Application detection types](#)” on page 70.

- 4 If this application detection method is the first for the policy, provide the required information in the dialog box, and click **OK**. If this application detection is an additional method, specify an operator for the detection method, provide the required information in the dialog box, and click **OK**.

See “[How operators work in application detection](#)” on page 72.

If multiple detection types are used, the **If** operator is only available for the first detection type.

- 5 Continue adding application detection methods until the application detection configuration is complete.

- 6 Click **Save changes**.

Application detection types

Application detection searches a monitored computer to determine if specific applications are installed.

See “[About application detection](#)” on page 68.

Application detection is configured for the monitor policies that are part of monitor packs . You can also add detection methods to a monitor policy to meet any additional application detection needs.

See “[Adding application detection to a monitor policy](#)” on page 69.

You can choose from several application detection types. The different types let you monitor the specific applications that are the most critical to your operations.

Table 5-1 Detection types

Type	Description
COM object is registered	(Windows only) Checks for the existence of a COM object. The Class ID and IID are required to identify the object. If you select the Always detect as a 32-bit module check box, applications are only detected if they run in a 32-bit state.

Table 5-1 Detection types (*continued*)

Type	Description
DLL is present	(Windows only) Checks for the presence of a DLL. The name of the DLL is required. If you select the Always detect as a 32-bit module check box, applications are only detected if they run in a 32-bit state.
File exists	Checks for the existence of a file. The path to the file and the name of the file are required. In Windows, you can check for the existence of a specific version or versions. To check for the existence of the file only, select Exists in the Condition field. To check for a specific product version or versions, select a logical operator and specify a version. Example: The version that is checked for can be equal to 7.0. In Linux or UNIX, you can check only for the existence of the file. Select Exists in the Condition field.
Package is present	Checks whether a package is installed on the monitored computer. For Linux systems, RPM package availability is checked. The package name is required. You can check for the existence of the package or a version of the package. To check for the existence of the package only, select Exists in the Condition field. To check for a specific version, select a logical operator and specify a version.
Process is running	Checks if a process is running. The name of the process is required. In Linux and UNIX environments, use the long name of the process without path or arguments. For example, for the "/usr/lib/dmi/snmpXdmid -s myhost" process, enter "snmpXdmid" in this field.
Registry key exists	(Windows only) Checks for the existence of a registry key. The registry key root and subkey are required. If you select the Always detect as a 32-bit module check box, applications are only detected if they run in a 32-bit state.
Registry value	(Windows only) Checks for the existence of a registry key value. The registry key root, subkey , key name, and key type are required. You can check for the existence of the value or the existence of a certain key value. To check for the existence of the value only, select Exists in the Condition field. To check for a specific key value, select a logical operator and specify a value. If you select the Always detect as a 32-bit module check box, applications are only detected if they run in a 32-bit state.

Table 5-1 Detection types (*continued*)

Type	Description
Service is installed or running	(Windows only) Checks for the existence or running of a service. The name of the service, and whether the name is a display or a binary name, is required. You must also choose if you want to check for the existence or running of the service.

How operators work in application detection

Operators determine the detection logic for the specified detection items. The detection logic tells Monitor Solution how to proceed through the detection methods to check for the application. The operators let you create more complex methods and allow for more accurate application detection. With only one detection method specified, your search may be more generic; the operators allow for a more focused search.

The detection methods that you specify are evaluated linearly—the functions that you specify evaluate one after the other. For example, if you set up methods stating the following:

IF a certain DLL is present

AND a certain file is present

OR a specific process is running

The detection logic evaluation first checks to see if the DLL and the file are present. If those are both present, then detection passes. If those are not present, then the logic moves on to check if the process is running. If the process is running, detection passes. If it is not running, detection fails. The set does not check first to see if the DLL is present, and then check to see if the file is present or the process is running. The "and" applies to the first step in the detection method.

If multiple detection types are used, the **If** operator is only available for the first detection type.

See “[Adding application detection to a monitor policy](#)” on page 69.

Adding computers to a monitor policy

After you create and configure a monitor policy you can apply the policy to groups of computers. After a monitor policy is saved and enabled, the policy is applied to the targeted computers on the next policy refresh. Policies refresh every one hour by default.

To add a computer to a monitor policy**1 Create or edit a monitor policy.**

To create a monitor policy See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63.

To edit an agent-based monitor policy See “[Editing agent-based monitor policies](#)” on page 66.

To edit an agentless monitor policy See “[Editing agentless monitor policies](#)” on page 67.

2 Click **Apply to and select one of the following options:**

- Select **Quick apply** to choose a list of common groups and targets, select a target, and click **Apply**.
- Select **Computers** to open the **Select computers** dialog box to search for groups of computers. Use the tool to locate computers, select the computers, and then click **OK**.
- (Agentless monitor policies only). Select **Resources** to open the **Select resources** dialog box to search for groups of computers and devices. Use the tool to locate resources, select the resources, and then click **OK**.

3 Click **Save changes.**

Adding computers to a monitor policy

Working with Metrics

This chapter includes the following topics:

- [About metrics](#)
- [Creating and editing metrics in the metric library](#)
- [Adding a metric to a rule](#)
- [Types of metrics](#)
- [About multiple instance metrics](#)

About metrics

Metrics are measurements of the data Monitor Solution collects from the computers and devices that you monitor. Metrics are used within rules to pinpoint problems and define their cause. Metrics define the types of data that you collect as well as the type of data source that you collect from. Metric data can be log events, the status of a program, or the status of an operating system component. Metric sources are the aspects of a computer or application that provide data.

When Monitor Solution and monitor packs are installed they include numerous predefined metrics for you to use. You can also create your own custom metrics.

See “[Creating and editing metrics in the metric library](#)” on page 75.

See “[Types of metrics](#)” on page 77.

Creating and editing metrics in the metric library

You can view, filter, create, edit, clone, and delete metrics from the metric library. Metrics are sort-able by metric type as well as by reference count. Reference count shows how many rules use the metric.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

To create and edit metrics in the metric library

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Metric Library**.

The metric library is divided into two sections: a section for agent-based metrics and a section for agentless metrics. You can expand or collapse the display of either section. In either section you can do any of the following actions:

Create and delete metrics You can click the **New** option in the toolbar and then select a metric type to create.

See “[Types of metrics](#)” on page 77.

You can also select a metric and click the **Delete** option in the toolbar to delete a metric. However, you cannot delete a metric that other rules or policies reference. Verify that a “0” is displayed in the **Referenced Count** column of the metric library to make sure that a metric is not referenced.

Edit existing metrics You can select a metric and then click the **Edit** option in the toolbar.

Warning: If you edit a metric that is referenced in by a rule, the rule is updated to include the updated metric. If you do not want this behavior to occur, first create a clone of the metric and then edit the clone.

Clone existing metrics You can select a metric and then click the **Clone** option in the toolbar. A clone of the metric is created in the library with **Copy of** prepended to the original name.

View and Search for metrics You can type a search string in the **Search** field and press Enter to display only the metrics that match your search criteria. You can also select a metric type from the view drop-down list to only display metrics of a certain metric type.

Adding a metric to a rule

Metrics return the data that the rules need to evaluate a condition and determine the necessary actions to take.

To add a metric to a rule

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Rule Library**.
- 3 In the right pane, double-click the rule to which you want to add a metric.
- 4 In the dialog box that opens, under **Metrics**, click the **New** symbol.
- 5 In the dialog box that opens, specify the metric parameters, and then click **OK**.
See “[About metric evaluation](#)” on page 111.
- 6 Click **OK**.

Types of metrics

This topic describes the different agent-based and agentless metric types. Agent-based metrics require that the Monitor Plug-in is installed on the target computer. Agentless metric do not require the Monitor Plug-in to be installed on the target computer.

The following are the different metric types:

Agent-based metric types	Agent-based metric types require that the Monitor Plug-in is installed on targeted computers. See Table 6-1 on page 78.
Agentless metric types	Agentless metric types do not require the Monitor Plug-in to be installed on targeted computers. See Table 6-2 on page 79.
Agent-based or agentless metric types	Agent-based and agentless metric types can be used with or without the Monitor Plug-in being installed on targeted computers. See Table 6-3 on page 79.

Table 6-1 Agent-based metric types

Metric type	Description
COM	Uses a custom COM DLL or other components to perform a custom data query. See “ New COM Metric page ” on page 80.
Command	Parses the results of a command-line utility or an input file for a particular value, and returns the associated values. See “ New Command Metric page ” on page 80.
Compound	Takes the data that one or more existing metrics collect and manipulates the values. See “ New Compound Metric page ” on page 84.
Custom DLL	Provides a way to perform a custom data query. See “ New Custom DLL Metric page ” on page 85.
Log event	Monitors the log files. See “ New Log Event Metric page ” on page 89.
Performance Counter	Collects data from Windows performance counters. Performance counters collect data such as CPU utilization, memory usage, disk swapping time, and processor cache usage. See “ New Performance Counter Metric page ” on page 92.
SQL	Retrieves data from Microsoft SQL Server databases on Microsoft Windows computers. The SQL Server database must be on the computer with the Monitor Plug-in. See “ New SQL Metric page ” on page 96.
Windows Process	Monitors currently running Windows processes. See “ New Windows Process Metric page ” on page 98.
Windows Service	Monitors currently running Windows services. See “ New Windows Service Metric page ” on page 99.

Table 6-2 Agentless metric types

Metric Type	Description
Group	Aggregates or groups data from several metrics into a single metric. See “ New Group Metric page ” on page 86.
HTTP	Remotely monitors a Web (HTTP) server's health. See “ New HTTP Metric page ” on page 87.
Ping	Monitors the response time of remote IP devices. See “ New Ping Metric page ” on page 91.
Smart	Automatically chooses the correct protocol and the correct authentication method to capture the metric data. See “ New Smart Metric page ” on page 96.
WS-MAN	Monitors Web Services for Management (WS-MAN) properties. See “ New WS-MAN Metric page ” on page 102.

Table 6-3 Agent-based and agentless metric types

Metric type	Description
Port	Monitors the computer's communication ports and makes sure that they are available for use. Ports are where connections can be initiated, maintained, and terminated. Data transfer occurs through ports, and services on computers respond to connections that come through them. See “ New Port Metric page ” on page 93.
SNMP	Polls attribute values from SNMP-enabled agents. See “ New SNMP Metric page ” on page 94.
WMI	Monitors various Windows Management Instrumentation (WMI) numeric properties. See “ New WMI Metric page ” on page 100.

New COM Metric page

This page lets you create and edit COM metrics. A COM metric uses a COM DLL to perform a data query on Windows computers. The query can be anything the DLL creator chooses based on the implementation of the DLL.

The COM metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-4 Options on the New COM Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Data type	The type of data that the metric retrieves: numeric or string.
Class ID	The Class ID of the COM object that provides the metric data.
Query string	The string that is passed to the COM object for the data that represents a custom data query.

New Command Metric page

This page lets you create and edit Command metrics. The command metric lets you parse the results of a command-line utility or an input file for a particular value. It then returns the associated values.

For example, you can run a command-line utility to list the names and creation dates of files in a specific location on a hard drive. This list can be parsed for the presence of a specific file name. If the file name is found, the creation date of the

file is returned. This example can be used to ensure that the proper version of an application is used.

The Command metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-5 Options on the New Command Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Data type	The type of data that the metric retrieves: numeric or string. If the data type selected is numeric and the input contains non-numeric values, the non-numeric values are converted to zeros (0).

Table 6-5 Options on the New Command Metric page (*continued*)

Option	Description
Use: Input file / Command line	<ul style="list-style-type: none"> ■ Input file – The file to parse. By default, you can process a maximum of 1000 lines in the input file. If you need to process more than this number, you need to set the parse line filter to specify the lines of data to process. For example, if you specify a range of 0-1004, the maximum number of lines that can be processed is 1005. ■ Command line – The command line to run, including parameters. The command utility or file must provide standard captured output. The file that is used or the command that is run must be on the local computer. To execute a command from the Windows shell (cmd.exe) you must specify a command line that includes cmd.exe. Then use the /c switch, to force the command to run and then to stop, and then the command parameter. For example: The time /t command must be entered as cmd /c time /t.
Line: Parsed lines / Skipped lines	<ul style="list-style-type: none"> ■ Parsed lines – specific lines to parse. All other lines are skipped. ■ Skipped lines – specific lines to skip. All other lines are parsed. <p>If these fields are left empty, all lines are parsed. The line numbering starts with 0.</p> <p>To specify more than one line or group of lines, separate the numbers with a comma.</p> <p>Lines can be specified in the following formats:</p> <ul style="list-style-type: none"> ■ X, Y – lines X and Y ■ First:XXX – first XXX lines ■ Last:XXX – last XXX lines ■ X-Y – lines X through Y inclusive

Table 6-5 Options on the New Command Metric page (*continued*)

Option	Description
Reverse parsing order	Reverses the data parsing order to read from right-to-left. The last column that is parsed is the left column. The column number order is also reversed. The right-most column is 0 and the column numbers are incremented going from right-to-left. This setting is arbitrary when you reference columns by name. The line numbers stay the same, so the top line of the file is still line 0.
Search pattern	Use this section to add and remove search patterns. A search pattern is a collection of conditions based on column values. You must specify a column to search and a value to search for. Rows are selected in accordance with the search pattern, and only these rows are processed. Other rows are ignored. After selection, the whole row is processed, not only the columns that were used in the search pattern.
Instance column	<p>Specifies the column by name or by column number value. If you use column names, the first row of the parsed data must contain the column names.</p> <p>Note: If you use a column name and specify lines to parse or skip, the first line that is returned is treated as a column name. Because the first line is treated as a column name, its value is not evaluated.</p> <p>If you use column numbers, the first column is on the left and starts at 0. The number's value is used as a label for the value that the Return value column returns.</p>
Return value column	<p>Specifies the column by name or the number in which the return value is located.</p> <p>Use the following guidelines:</p> <ul style="list-style-type: none"> ■ For column numbers use, the first column is on the left and starts at 0. ■ For column name use, the first row of the parsed data must contain the column names. <p>Note: If you use a column name and specify lines to parse or skip, the first line that is returned is treated as a column name. Because the first line is treated as a column name, its value is not evaluated.</p>
Column delimiter	Specifies custom delimiter to parse to the column.

New Compound Metric page

This page lets you create and edit Compound metrics. A compound metric takes the data that one or more existing metrics collect, and it arithmetically manipulates the data values. For example, you can use a compound metric to calculate the average value between two performance counter metrics.

Metrics that a compound metric references are polled synchronously at the time the compound metric is polled. This polling is in addition to the polling interval that is defined in the metric.

Compound metrics support the following arithmetic functions: +, -, *, /, and %

Note: A compound metric cannot use another compound metric.

The Compound metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-6 Options on the New Compound Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Metric variables	The list of the metrics that are used in the compound metric, along with their assigned variable name. The Add option lets you add metrics to the list.
Equation	The equation that defines the compound metric. You can type variable names and operators into the field. You can also insert the operators by selecting the operator in the list and clicking the Add option. Variable names are case-sensitive.

Table 6-6 Options on the New Compound Metric page (*continued*)

Option	Description
Use multiple instances	If the multiple instances option is specified, the following options are available: <ul style="list-style-type: none">■ Each metric used in the compound metric must contain the same number of instances. The evaluation of the compound metric is run one time for each instance of each metric that is used.■ One metric returns a single value and another returns multiple instances. The single value metric is used with each of the instance values of the multiple instance metric. If the metric does not comply with one of these options, it fails. See “ About multiple instance metrics ” on page 104.

New Custom DLL Metric page

This page lets you create and edit Custom DLL metrics. A custom DLL provides a way to perform a custom data query.

- In Windows systems, a custom DLL metric is similar to a custom COM object metric. The query can be anything the DLL creator chooses based on the implementation.
- In Linux and UNIX, instead of a custom DLL, the functionality resides in a UNIX shared library.

When using the custom DLL metric in Linux with the GCC compiler, link it with the `-Bsymbolic` flag. For example, `gcc -fPIC -shared -Wl,-Bsymbolic dlltest.o -o DllTest.so`

Note: The shared library should be compiled on a 32-bit platform.

The Custom DLL metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-7 Options on the New Custom DLL Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Data type	The type of data that the metric retrieves: numeric or string. If the data type selected is numeric and the input contains non-numeric values, the non-numeric values are converted to zeros (0).
DLL pathname	The path to the DLL, including the name of the DLL file.
Query string	A string that is passed to the DLL that represents a custom data query.

New Group Metric page

This page lets you create a group metric. A group metric aggregates the data a single metric collects across multiple resources. The data comes from a pre-selected set of resources that is referred to as the group. A group is the same as a policy target. A group can be a predefined list of resources, such as **Network Resources**, **All Printers**, and **All computers without installed Software Management Plug-in**. A group can also be a list of resources that are created by the user.

The metric you select to collect the data must be an agentless metric, since agent-based metrics only have a single resource.

When you create a group metric, you specify the operation to perform on the data. The operation options are minimum, maximum, average, and sum.

The Group metric does not use the Monitor Plug-in.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-8 Options on the New Group Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Thread Pool	Controls which of three thread pools is used to poll the metric source. Each thread pool's settings are configured in the Remote Monitoring Server Settings page. See " Remote Monitoring Server Settings: Performance Tuning tab " on page 57.
Metric	Specifies the metric that collects data from each of the selected resources. The metric must be agentless.
Operation	Specifies the operation that is performed on the collected data. The options are Max , Min , Sum , and Average .
Apply to	Specifies the computer or resource to which the metric applies. The Quick apply option lets you type in the group, filter, or target resource.

New HTTP Metric page

This page lets you create and edit HTTP metrics. The HTTP metric lets you remotely monitor the health of a Web (HTTP) server.

The HTTP metric does not use the Monitor Plug-in.

See "[About metrics](#)" on page 75.

See "[Types of metrics](#)" on page 77.

See "[Creating and editing metrics in the metric library](#)" on page 75.

Table 6-9 Options on the New HTTP Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.

Table 6-9 Options on the New HTTP Metric page (*continued*)

Option	Description
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Thread Pool	Controls which of three thread pools that the metric poll through. Each thread pool's settings are configured on the Remote Monitoring Server Settings page. See " "Remote Monitoring Server Settings: Performance Tuning tab" " on page 57.
SSL	Enables or disables SSL encryption.
Port number	The port number to use. When SSL is enabled, port 443 is used by default. When SSL is not enabled, port 80 is used.
Path	The path to the targeted host. This path is the relative path to the IP of the resource.

Table 6-9 Options on the New HTTP Metric page (*continued*)

Option	Description
Return value	<p>The value that the metric returns.</p> <p>The options are as follows:</p> <ul style="list-style-type: none">■ Content changed Initiates a CRC process on the HTML page that is retrieved. The value is stored. Future polls then evaluate the CRC against the value and return if the content of the page has changed or not. Note: This setting is only intended for polling static HTML pages. This option is not recommended for polling dynamic HTML pages as they can potentially change with every poll.■ Load time The amount of time that passes between when the first packet is sent and the last packet is received.■ Page size The file size of the page that is returned.■ Response time The amount of time that passes between when the first packet is sent and the first packet is received.■ HTTP status If a page is either successfully returned or fails.

New Log Event Metric page

This page lets you create and edit Log Event metrics. For the most common log types and formats, the necessary configuration is already done. For example, no special configuration is necessary to collect data from a Microsoft IIS Web log. If your log type or format is not predefined, you can specify the configuration settings.

Note: For log files that are moved and re-created, the metric's interval must be less than the time that is required to fill a log file. Move the file to a new location, create a new log file with the same name.

Warning: To monitor IIS logs, you must enable the extended logging property in the Internet Information Services utility on the monitored computers. If extended-logging is not enabled, Monitor Solution cannot properly monitor IIS logs. Rules can trigger falsely when this setting is not enabled.

The Log Event metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-10 Options on the New Log Event Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Log type	The type of log to monitor. Only the applicable options appear.
Custom location	If you select the Custom option for the Log type field, then the Custom location field is available. When using a custom location, enter the path of the folder that contains the log files or enter the path and the name of the log file. If you enter the path of the folder that contains the log files, all log files in that folder are monitored. The log file must be on the local computer; you cannot monitor logs on different computers.
Log format	The format of the log to monitor.
Custom format	This option is for entering custom column names. Separate the column names with specified delimiters.
Custom delimiter	For entering delimiters when you monitor custom log formats.

New Ping Metric page

This page lets you create and edit Ping metrics. The Ping metric type lets you remotely monitor the response time of remote IP devices.

The Ping metric does not use the Monitor Plug-in.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-11 Options on the New Ping Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Thread Pool	(Agentless only) Controls which of three thread pools that the metric poll will use. Each thread pool’s settings are configured in the Remote Monitoring Server Settings page. See “ Remote Monitoring Server Settings: Performance Tuning tab ” on page 57.
Number of packets	The number of packets sent. The ping metric is declared a success only when all packets that are sent are returned. If your environment requires a very low threshold for packets dropped, you can increase the number of packets sent. If your environment allows for a higher threshold of dropped packets, you can decrease the number of packets sent.
Packet size	Adjusts the packet size to either 32 bytes or 64 bytes. Packet size can be adjusted to simulate different types of traffic in the network.

Table 6-11 Options on the New Ping Metric page (*continued*)

Option	Description
Return value	<p>The data that the metric returns.</p> <p>The options are as follows:</p> <ul style="list-style-type: none"> ■ Number of hops Returns the number of hops of a trace route. ■ Number of packets Returns how many packets were successfully returned. ■ Ping status Returns that the ping is a success or a failure. To be successful, all packets must be returned. ■ Return trip time Returns the minimum round-trip time, the maximum round-trip time, and the average round-trip time.
All instances	If this option is checked, the rule is triggered only if all instances meet the required conditions. If this option is not checked, the rule is triggered if any instance meets the required conditions.

New Performance Counter Metric page

This page lets you create and edit Performance Counter metrics. Windows operating system performance counters collect data on the performance of the computer. The Monitor Plug-in can collect data from these counters to monitor numerous performance aspects. Performance aspects can include CPU utilization, memory usage, disk swapping time, and processor cache usage.

The Performance Counter metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-12 Options on the New Performance Counter Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.

Table 6-12 Options on the New Performance Counter Metric page (*continued*)

Option	Description
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Performance Counter Builder	Opens the Performance Counter Builder dialog box, which helps you create the performance counter metric. The applicable values are automatically inserted into the Performance object, Counter, and Instance fields. Note: The values are automatically taken from the Notification Server computer and may not exist on the computer that the metric monitors.
Performance object	The performance counter object name.
Counter	The performance counter name.
Instances	The performance counter instance.
All instances	This option forces a metric to return all instances of the metric data. See " About multiple instance metrics " on page 104.

New Port Metric page

This page lets you create and edit Port metrics. The Port metric lets you monitor a port to ensure that the port is available for communications. Ports are places where connections can be initiated, maintained, and terminated. Data transfer occurs through ports, and services on a computer respond to connections that come through ports.

The Port metric can be used to monitor computers whether or not the Monitor Plug-in is installed.

See "[About metrics](#)" on page 75.

See "[Types of metrics](#)" on page 77.

See "[Creating and editing metrics in the metric library](#)" on page 75.

Table 6-13 Options on the New Port Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Thread Pool	(Agentless only) Controls which of three thread pools that the metric poll will by. Each thread pool's settings are configured in the Remote Monitoring Server Settings page. See “Remote Monitoring Server Settings: Performance Tuning tab” on page 57.
Port number	The port number to be monitored.

New SNMP Metric page

This page lets you create and edit SNMP metrics. The SNMP metric polls attribute values from SNMP enabled Monitor Plug-ins. SNMP uses a distributed architecture that consists of management applications and agent applications. In Monitor Solution, the management application is the Symantec Management Console, which is used for network management functions.

Note: Although SNMP is a TCP/IP protocol, the Monitor Plug-in can only get SNMP attribute values from the computer on which it is installed.

Note: On all supported systems, you must configure security for SNMP metrics to start working. By default, no communities are added to the “accepted community names” list.

The SNMP metric can be used to monitor computers whether or not the Monitor Plug-in is installed.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-14 Options on the New SNMP metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	(Agent-based only) The amount of time in seconds to wait for a response before the metric data collection fails.
Data type	The type of data that the metric retrieves: numeric or string. If the data type selected is numeric and the input contains non-numeric values, the non-numeric values are converted to zeros (0).
Thread Pool	(Agentless only) Controls which of three thread pools that the metric poll will by. Each thread pool’s settings are configured in the Remote Monitoring Server Settings page. See “ Remote Monitoring Server Settings: Performance Tuning tab ” on page 57.
SNMP OID	The object identifier, which uniquely identifies the object variable in the Management Information Bases (MIB). The OID must reference either a scalar attribute or a member attribute of a table entry. OIDs that reference the entire table entry are not valid. Agent-based OID can contain OIDs like <code>1.2.3.65.4.</code> , however agentless OIDs should be like <code>.1.2.334.5.4.3</code> .
Use multiple instances	This option allows the metric to retrieve multiple instances of the data. See “ About multiple instance metrics ” on page 104.

New Smart Metric page

This page lets you create and edit Smart metrics. Smart metrics remotely collect metric data from computers. Smart metrics automatically choose the correct protocol and the correct authentication method to capture the metric data. The logic is based on the information that is gathered during a network discovery and is stored in the Protocol Abstraction Layer.

The Smart metric does not use the Monitor Plug-in.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-15 Options on the New Smart Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	<p>The amount of time in seconds between when the metric source is polled for data.</p> <p>This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.</p>
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Thread Pool	<p>Controls which of three thread pools that the metric poll will by. Each thread pool’s settings are configured in the Remote Monitoring Server Settings page.</p> <p>See “Remote Monitoring Server Settings: Performance Tuning tab” on page 57.</p>
Smart Key	Opens the Select Smart Key dialog box where you select the metric data that is monitored.

New SQL Metric page

This page lets you create and edit SQL metrics. The SQL metric lets you retrieve data from Microsoft SQL Server databases on Windows computers.

The SQL metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-16 Options on the New SQL metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Recordmetric values every field of the Monitor Plug-in configuration.
Data type	The type of data that the metric retrieves: numeric or string. If the data type selected is numeric and the input contains non-numeric values, the non-numeric values are converted to zeros (0).
Override Monitor Plug-in settings	If you select Override Monitor Plug-in settings, then you must specify a Server name and a Database name to connect to. You must also specify one of the following connection methods: <ul style="list-style-type: none">■ Windows Authentication You can use this field to specify a Windows account User name and Password to use to authenticate with.■ SQL Authentication You can use this field to specify the SQL Server User name and Password to authenticate with.■ NS Authentication If this option is selected the credentials for the Notification Server computer are used to authenticate.

Table 6-16 Options on the New SQL metric page (*continued*)

Option	Description
SQL Query Builder	Assists you to create a query. You can select the database table, the column to use and a condition. Note: The SQL Query Builder uses the authentication settings on the Plug-in Configuration page to connect to the SQL Server database. If you receive an error when you open the SQL Query Builder , make sure that the Plug-in Configuration authorization settings match the Notification Server computer SQL Server authentication settings.
SQL query	The SQL query to use to retrieve the data.
Instance column	The column in the database for retrieving instance names. Multiple instance metrics use this value to distinguish between the different instances. See “ About multiple instance metrics ” on page 104.
Return value column	The name of the database’s return value column.
Use multiple instances	Select this check box to allow the metric to retrieve multiple instances of the data. See “ About multiple instance metrics ” on page 104.

New Windows Process Metric page

This page lets you create and edit Windows Process metrics. The Windows Process metric type lets you determine whether-or-not a specific process is running on a Windows 32-bit computer.

The Windows Process metric requires the Monitor Plug-in to be installed on the targeted computer.

The Windows Process metric returns the count of process instances that are running. For example, if there are two processes that are running with the same, then the metric returns "2" as the result.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-17 Options on the New Windows Process Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Process name	The name of the Windows Process that is monitored.

New Windows Service Metric page

This page lets you create and edit Windows Service metrics. The Service metric type lets you monitor services on Windows 32-bit computers.

You can use Windows Service metrics to determine the following about a service:

- Installed
- Running
- Enabled or disabled on computer startup

The Windows Service metric requires the Monitor Plug-in to be installed on the targeted computer.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-18 Options on the New Windows Service Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.

Table 6-18 Options on the New Windows Service Metric page (*continued*)

Option	Description
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in's metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	The amount of time in seconds to wait for a response before the metric data collection fails.
Service name	The name of the Windows service that is monitored.
Query type	The information that you require about the service. The options are as follows: <ul style="list-style-type: none"> ■ Installed Determines if the service is installed on the computer or not. ■ Status Determines if the service is running on a computer or not. ■ Startup Type Determines if the service is enabled or disabled when the computer starts-up.

New WMI Metric page

(Windows only)

This page lets you create and edit WMI metrics. Monitor Solution can monitor various Windows Management Instrumentation (WMI) numeric properties. WMI includes namespaces, classes, and instances. The WMI namespace is a unit for grouping classes and instances. A class can be thought of as a unit of management. In WMI, a hardware or a software system component is called an object and is represented as an instance of a WMI class. An instance is the representation of a managed object. The property of the object instance contains the numeric value that you want to monitor.

You can identify WMI numeric properties to track. For example, Monitor Solution can monitor the free disk space on a computer by collecting a value at a specified interval. You can use this metric data for capacity planning. Monitoring the free

disk space provides useful information for you to use when you plan for upgrades or additional drives.

Monitor Solution can also monitor computers for performance thresholds using WMI. You can create a threshold value, and when the threshold is crossed, Monitor Solution triggers a rule. For example, you can set a threshold value for free disk space at greater than 85%. If the sampled free disk space value is greater than the 85% threshold, then Monitor Solution can trigger a rule.

WMI metrics that poll too frequently (every 60 seconds or less) may cause significant CPU usage on the monitored computer.

The WMI metric can be used to monitor computers whether or not the Monitor Plug-in is installed.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-19 Options on the New WMI Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Timeout	(Agent-based only) The amount of time in seconds to wait for a response before the metric data collection fails.
Data type	The type of data that the metric retrieves: numeric or string. If the data type selected is numeric and the input contains non-numeric values, the non-numeric values are converted to zeros (0).
Thread Pool	(Agentless only) Controls which of three thread pools that the metric poll will by. Each thread pool’s settings are configured on the Remote Monitoring Server Settings page. See “ Remote Monitoring Server Settings: Performance Tuning tab ” on page 57.

Table 6-19 Options on the New WMI Metric page (*continued*)

Option	Description
Namespace	The WMI Namespace for the object path. For example, root\default. For more information, see the WMI SDK.
Property	Within WMI namespaces there are various WMI classes (for example: WIN32_Process). Each class has its own properties and methods. For example, a property of the WIN32_Process class is "HandleCount". A method of this class is "GetOwner". If you choose the Property option, you must specify the class name of a WMI class and the counter name to a property. The WMI properties are as follows: <ul style="list-style-type: none">■ Class name The class name for the object path. This class name is required when you don't use a WQL string (for example: win32_useraccount).■ Counter The Property name for the object. This field is required even when used with a WQL string (for example: accounttype).■ Instance The Instance information for the object path. This field is required when you don't use a WQL string (for example: name="administrator", domain="Symantec"). Note: All keys that define an instance must be specified. Multiple keys are comma-delimited. The order of keys is arbitrary and the keys' names and values are not case sensitive.
Query	Specifies that a WMI Query Language string is used instead of a class name.

New WS-MAN Metric page

This page lets you use Web Services for Management (WS-MAN) to poll data for metrics. WS-MAN provides a common way for systems to access management information and exchange management information across the IT infrastructure. This protocol is similar in its design to Web-based Enterprise Management (WBEM),

which is implemented in Windows as Windows Management Instrumentation (WMI). Like WMI, WS-MAN can be used to get or set variables. It can also be used to perform more complex operations, such as method execution.

Monitor Solution uses WS-MAN like it uses WMI. First, a namespace is given, to which Monitor Solution tries to connect. In addition to WMI namespaces, a wider variety of WS-MAN namespaces can be used. Returned data can be either numeric or string, depending on the type of the attribute requested.

The WS-MAN metric does not use the Monitor Plug-in.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

Table 6-20 Options on the New WS-MAN Metric page

Option	Description
Name and Description	A name and description to help you identify and locate the metric.
Polling interval	The amount of time in seconds between when the metric source is polled for data. This interval should not be set to less than the Monitor Plug-in’s metric collection interval. The metric collection interval is specified in the Record metric values every field of the Monitor Plug-in configuration.
Data type	The type of data that the metric retrieves: numeric or string. If the data type selected is numeric and the input contains non-numeric values, the non-numeric values are converted to zeros (0).
Thread pool	Controls which of three thread pools that the metric poll through. Each thread pool’s settings are configured on the Remote Monitoring Server Settings page. See “ Remote Monitoring Server Settings: Performance Tuning tab ” on page 57.
Namespace	Specifies the namespace to which the metric connects to poll data.

Table 6-20 Options on the New WS-MAN Metric page (*continued*)

Option	Description
Property	Lets you specify a certain attribute from a certain class instance. When you select this option, specify the following values: <ul style="list-style-type: none">■ Class name - Class from which to query (for example, CIM_OperatingSystem).■ Counter - Attribute of the class to retrieve (for example, Version).■ Instance - Specific instance of the class from which to query.
Query	Lets you provide a WQL string to be used to find the results. Selecting this option lets you specify a WQL String . A WQL String is used to execute a WQL query (for example, SELECT Description FROM CIM_ComputerSystem). If you specify a CIM_Query in an agentless WS-MAN Metric, you cannot receive any requested data during monitoring of the ESX Server. Use of the CIM_Query is followed with a warning in Log Viewer that states following actions are not supported by the service.

About multiple instance metrics

A multiple instance metric is a metric that can return more than one value in a query. For example, if you want to know disk space usage, there is a disk usage value available for each of the disks on the computer. With a metric specified as a multiple instance metric, you can receive values for each of the disks. If the metric is not configured for multiple instances, then a single value is retrieved for the first instance of the query.

Metrics are configured to use multiple instances by checking the **Use multiple instances** option when you create or edit a metric.

See “[About metrics](#)” on page 75.

See “[Types of metrics](#)” on page 77.

See “[Creating and editing metrics in the metric library](#)” on page 75.

The metric types that support multiple instances are as follows:

- Compound

See “[New Compound Metric page](#)” on page 84.

- Performance Counter
See “[New Performance Counter Metric page](#)” on page 92.
- SNMP
See “[New SNMP Metric page](#)” on page 94.
- SQL
See “[New SQL Metric page](#)” on page 96.

Working with rules

This chapter includes the following topics:

- [About rules](#)
- [Creating and editing rules in the rule library](#)
- [Cloning rules](#)
- [Creating and editing rules](#)
- [About metric evaluation](#)
- [New Metric Evaluation page for log event rules](#)
- [New Metric Evaluation page for Metric rule types](#)
- [New Metric Evaluation page for NT Event rule types](#)
- [Types of rules](#)

About rules

Rules are used within monitor policies to specify what metrics to monitor and how fluctuations in a metric's data value should be interpreted. After a rule is defined and activated, the metric data is collected from a specified application or operating system.

Rules define the following information:

- The metrics that are used to collect data from metric sources
- The acceptable values for metric data
- When the rules are triggered
- The actions that are taken when metric data is in an undesired state or is beyond a wanted range

See “[About metrics](#)” on page 75.

Monitor Solution has different types of rules for different purposes. Rules can be used to poll metric sources, such as Windows NT events and log events. Some rules are configured to only collect metric data and nothing else. Other rules both collect data and also evaluate it against predefined acceptable ranges and values. If a data value is in an undesirable state or is beyond a wanted range, then the rule is considered triggered. When a rule triggers, actions can be taken.

See “[Types of rules](#)” on page 115.

The actions that a rule can trigger are defined either within the rule or within the monitor policy. Actions, such as running a script, can be initiated from the monitored computer. Actions, such as sending an email, can also be run as server-side tasks from the Notification Server computer.

See “[About Monitor Solution tasks and actions](#)” on page 117.

Creating and editing rules in the rule library

The rule library contains all of your rules. You can view and filter, create, edit, clone, and delete rules from the rule library. Rules are sort-able by type as well as by reference count. Reference count shows how many policies use the metric. Multiple monitor policies can reference a single rule. If a rule is edited, it is updated in all of the monitor policies that reference it.

To work with rules in the rule library

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Rule Library**.

The rule library is divided into two sections: a section for agent-based rule types and a section for agentless rule types. You can expand or collapse the display of either section.

In either section you can do any of the following actions:

Create and delete rules

You can click the **New** option in the toolbar and then select a rule type to create.

See “[Types of rules](#)” on page 115.

You can also select a rule and click the **Delete** option in the toolbar to delete a rule. However, you cannot delete a rule that other rules or policies reference.

Verify that a “0” is displayed in the **Reference Count** column of the rule library to make sure that a rule is not referenced.

See “[Creating and editing rules](#)” on page 110.

Edit an existing rule

You can select a rule and then click the **Edit** option in the toolbar.

See “[Creating and editing rules](#)” on page 110.

Warning: If you edit a rule that is referenced in a policy, the policy is updated to include the updated rule. If you do not want this behavior to occur, first create a clone of the rule and then edit the clone.

Clone an existing rule

You can select a rule and then click the **Clone** option in the toolbar. A clone of the rule is created in the library with **Copy of** prepended to the original name.

See “[Cloning rules](#)” on page 110.

View and Search for rules

You can type a search string in the **Search** field and press Enter to display only the rules that match your search criteria. You can also select a category type from the view drop-down menu to only display rules of a certain category.

Cloning rules

You can clone rules to create a new rule that is similar to an existing rule.

To clone a rule

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Rule Library**.
- 3 In either the **Agent-based** rules table or in the **Agentless** rules table, select the rule to clone.
- 4 In the toolbar, click the **Clone** symbol.

A clone of the rule is created in the library with **Copy of** prepended to the original name.

- 5 Select the newly created rule and edit it as needed.

See “[Creating and editing rules](#)” on page 110.

Creating and editing rules

You can create rules from the rule library.

See “[Creating and editing rules in the rule library](#)” on page 108.

To create a rule

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Policies > Rule Library**.
- 3 In either the **Agent-based** rules table or in the **Agentless** rules table, do one of the following.

Create a new rule

- In the toolbar, click the **New** symbol.
- Select a rule type.
See “[Types of rules](#)” on page 115.
- Enter a name and a description.
- Select a category.

Edit a rule

- Locate and select the rule that you want to edit.
- In the toolbar, click the **Edit** symbol.
- (Optional) Edit the name, description, and category for the rule.

Warning: If you edit a rule that is referenced in a policy, the policy is updated to include the updated rule. If you do not want this behavior to occur, first create a clone of the rule and then edit the clone.

- 4 If applicable to the rule type, under **Metrics**, add metric evaluation logic to the rule.

You can use the **Repeat count** and **Time period** fields to specify the number of times the metric must trigger in a specified period of time before the rule triggers.

See “[About metric evaluation](#)” on page 111.

- 5 If applicable to the rule type, under **Actions**, configure severity and task settings for the rule should it become triggered.

See “[About severity states](#)” on page 119.

See “[Adding tokens to a Send Email task](#)” on page 123.

- 6 Click **OK**.

About metric evaluation

Metric evaluation is the assessment of metric data. Metric evaluation lets you compare the metric data to a pre-determined value or range or the value of another metric. If the metric data meets the specified conditions of the rule, then the rule is triggered. Triggered rules can perform actions such as raising an alert in the Event Console.

See “[About Monitor Solution tasks and actions](#)” on page 117.

Rules are classified into different rule types. Some rule types use metric evaluation and some rule types do not. If a rule type supports metric evaluation, the rule can evaluate the metric data that it gathers. If a rule type does not support metric evaluation, the rule gathers the metric data but does not evaluate it.

See “[Types of rules](#)” on page 115.

Metric evaluation supports the following metric types:

- Log Event

See “[New Metric Evaluation page for log event rules](#)” on page 112.

- Metric

See “[New Metric Evaluation page for log event rules](#)” on page 112.

- NT Event

See “[New Metric Evaluation page for NT Event rule types](#)” on page 114.

You can add multiple sequences of metric evaluations to a rule using the IF, AND, and OR operators. The operators determine the logic for the evaluations. The operators let you create more complex rules and allow for more accurate evaluations. With only one metric evaluation, your rule may be more generic—the operators allow for a more focused rule.

The evaluation methods that you specify are processed linearly—the functions that you specify evaluate one after the other. For example, if you set up a rule that stated that:

IF metric evaluation x.

AND metric evaluation y.

OR metric evaluation z.

The evaluation first checks to see if both metric evaluation x and metric evaluation y are true. If those are both true then the rule is triggered. If not, then the logic moves on to check if metric evaluation z is true. If metric evaluation z is true then the rule is triggered. This set of evaluation methods does not check first to see if metric evaluation x is true, and then check to see if metric evaluation y is true or metric evaluation z is true. The "and" applies to the first step in the evaluation method.

Note: If multiple evaluations are used, the **If** operator is only available for the first evaluation.

New Metric Evaluation page for log event rules

This page lets you set up metric evaluation conditions for log event rule types.

See “[About metric evaluation](#)” on page 111.

Table 7-1 Options on the Metric evaluation page

Option	Description
Metric	The log event metric to use as the metric source.
Property	The log event property that is used as the metric source. This value is set to "all data" when you select Unformatted .

Table 7-1 Options on the Metric evaluation page (*continued*)

Option	Description
Condition	The condition that the log event value must meet.
Value	The log event value.

New Metric Evaluation page for Metric rule types

This page lets you set up metric evaluation conditions for Metric rule types.

See “[About metric evaluation](#)” on page 111.

Table 7-2 Options on the Metric Evaluation page

Option	Description
Metric	The list of metrics to evaluate.
Uses profiled data	Enables the plug-in profiling feature for the metrics that are associated with the rule.
Statistics	<p>Specifies the statistical function that is applied to the data. The data is retrieved from the metric source over the interval that is specified in the Time period field. The options depend on the type of metric that is selected in the Metric option.</p> <p>The options are defined as follows:</p> <ul style="list-style-type: none"> ■ Absolute deviation – The absolute deviation of the collected data. ■ Average – The average or mean of the collected data. ■ Maximum – The maximum value of the collected data. ■ Minimum – The minimum value of the collected data. ■ Rate – The difference of the final and the initial time period values, divided by the total time period value. ■ Standard deviation – The standard deviation of the collected data.
Time period	<p>The amount of time over which the data is collected, including the unit of time. This value is used in the statistical calculation. For example, if the value in the Statistics field is average and the value in this field is 10 seconds, the average of the data that is collected over 10 seconds is used.</p> <p>Set this value to be greater than the metric update interval.</p>

Table 7-2Options on the Metric Evaluation page (*continued*)

Option	Description
Condition	Specifies the comparison condition between the metric value or values (for multiple instance metrics) and the value in the Value field. If the metric value meets the test, the rule is triggered.
Value type	The type of value to which the associated metric is compared to determine if the rule needs to be triggered. The options are as follows: <ul style="list-style-type: none"> ■ Constant – You must enter a value with which you want the associated metric to be compared. ■ Metric – The associated metric is compared to the value of another metric that already exists.
Value	The metric value or constant to whose value the associated metric value is compared to determine if the rule needs to be triggered. The options for this field are the defined metrics of the type that is specified in the Value type field. If Constant is selected in the Value type field, enter a value to which the associated metric is compared. If a metric type is selected, enter the metric to use. If the selected metric contains multiple instances, each instance value is compared instance-to-instance, regardless of the setting of the Match multiple instance when evaluating rule check box. If the Use profiled data check box is selected, the threshold is used. The threshold sets a value relative to the average value for the associated metric in standard deviations.

New Metric Evaluation page for NT Event rule types

This page lets you set up metric evaluation conditions for NT event rule types.

See “[About metric evaluation](#)” on page 111.

Table 7-3 Options on the Metric evaluation page

Option	Description
Operator	Specifies how the current NT event property is logically joined to the previous NT event property. The options are “And” and “Or”. This field is not applicable to the first or to the only NT event property.
Property	The NT event property that is used as the metric source.
Condition	The condition that the NT Event property's value must meet.
Value	The value of the NT event property.

Types of rules

Monitor Solution uses rules within monitor policies to collect and evaluate metric data. Rules have six different types, called rule types. Each rule type is used for different purposes. Agent-based monitor policies support all rule types. Agentless monitor policies only support the **Metric** rule types and the **Metric Collect** rule types.

See “[About rules](#)” on page 107.

Table 7-4 Rules types

Rule type	Description
Metric	Uses polled metrics to gather metric data that is compared against a predetermined value. If the specified conditions of the metric data are reached then the rule becomes triggered. When rules are triggered, the severity state of the rule changes and the actions for that rule are run.
Metric Collect	Gathers metric data but does not evaluate it. Metric Collect rule types do not have any associated severity state or actions. Metric Collect is used if you want to collect data, but you do not need the rule to run actions.
NT Event	The rule type NT Event is based on Windows NT events. Whenever a Windows NT event occurs on the monitored computer, the event is evaluated against all the NT event rules. If any of the rules are triggered, the severity state of the rule changes and the actions for that rule are run.

Table 7-4 Rules types (*continued*)

Rule type	Description
NT Event Collect	Gathers NT Event data. The NT Event Collect is not evaluated but it does not have any associated severity state or actions. NT Event Collect is used if you want to collect NT Event data, but you do not need the rule to run actions.
Log Event	The rule type Log Event is based on log-events. Whenever a log-event triggers a rule, the severity state of the rule changes and the actions for that rule are run.

Working with tasks and actions

This chapter includes the following topics:

- [About Monitor Solution tasks and actions](#)
- [About severity states](#)
- [Monitor task types](#)
- [Adding tokens to a Send Email task](#)
- [Adding actions to rules](#)
- [Adding actions to monitor policies](#)
- [Monitor client and server token types](#)

About Monitor Solution tasks and actions

Monitor tasks and actions can be scheduled, run on demand, or run as part of rules and policies. You can run tasks from a task server or you can choose from several Monitor-specific task types.

See “[Monitor task types](#)” on page 120.

You can add tasks and actions to a rule or a policy. Rules are triggered when monitored metric data reaches a determined value or goes beyond an acceptable value range. A triggered rule sends an alert, and any actions or tasks that are specified for that rule are executed.

When actions are part of a policy, the point at which the actions are executed depends upon the severity setting of the action. When actions are associated with a rule, they are executed when that rule is triggered. When actions are associated

with a policy, they are assigned a severity state. The actions are executed when a rule with that same severity that is specified for that policy is triggered. Within a policy, each severity state can have an action or set of actions specified for it. For example, you can have rules specified for a policy that have a severity state of Critical. When any of the critical rules are triggered, all of the actions that are specified for the Critical severity state are executed.

Setting actions at a rule level makes the actions more specific by targeting an individual metric. Setting actions at a policy level lets you specify actions so that they are executed to respond to multiple sources. The disadvantage is that the actions may have to be more general.

You can include the same task in multiple rules or policies. Modifying a task in a rule or policy also changes that rule in any other rules or policies that use that task.

You can specify task server actions or Monitor Plug-in actions for your rules and policies. Task server actions are run from the task server, and Monitor Plug-in actions are run from the Monitor Plug-in. Agentless policies can only contain task server actions.

The advantages of using task server tasks are as follows:

- You can create jobs from the task server.
- You can easily get history information from task server tasks by viewing the task item.
- More tokens are available for configuring the tasks than there are for Monitor Plug-in tasks.
- More task types available than there are for Monitor Plug-in tasks.

The advantages and disadvantages of using Monitor Plug-in tasks are as follows:

- Tasks can be run even if the Notification Server computer is not reachable, which may make Monitor Plug-in tasks very useful for critical tasks.
- Not as many task types are available as there are for task server tasks.
- Not as many tokens are available for configuring the tasks as there are for task server tasks.
- You can only create client tasks.
- You cannot create jobs with Monitor Plug-in tasks.

Monitor packs include predefined tasks to meet many of your needs. You can customize tasks further so that the execution of tasks is more useful.

Customizing monitor tasks lets you add additional monitor information to the tasks that are run when a rule is triggered, or as part of a policy. Additionally, you

can run tasks independently of any rules or policies. Tasks can be run on demand or on a schedule. With the additional Monitor data, the execution of the tasks is more meaningful.

See “[Adding tokens to a Send Email task](#)” on page 123.

You can add actions to rules.

See “[Adding actions to rules](#)” on page 124.

You can add actions to policies.

See “[Adding actions to monitor policies](#)” on page 125.

About severity states

Each rule has a severity state that is associated with it. The severity state of a resource reflects the severity level of rules that have been triggered on that resource.

See “[About rules](#)” on page 107.

The available severity states, from least severe to most severe, are as follows:

- Normal
- Undetermined
- Informational
- Warning
- Major
- Critical

When a rule is triggered, the severity state of that rule is visible in the Event Console. The state of the resource is set to the most critical severity level of any triggered rule. For example, two rules can trigger on a resource, one with a severity level of Warning and one with a severity level of Major. In that case, the overall state of the resource is Major. Normal is the base severity state for a rule when a rule is not triggered.

When you select a severity setting for a rule, you also choose how the severity is reset for the rule. When a rule is triggered, the severity for that rule changes for that rule, and an alert is sent to the Event Console. When the metric data goes back to an acceptable value range, the rule needs to be reset.

You can choose to reset a severity level in one of the following ways:

- Manual

The alert needs to be manually resolved in the Event Console.

- Updated Metric

If the metric for the rule crosses back to acceptable levels, the alert is resolved automatically.

- Updated Rule

This option is only available for NT Event rules and Log Event Rules. These types of rules do not have a threshold, so updated metrics cannot reset them. You can set up these rules so that the triggering of another rule of the same type can reset them. For example, an NT Event rule can reset another NT Event rule.

You can also specify an action or group of actions that runs for each severity state. When actions are associated with a policy, they are assigned a severity state. The actions are executed when a rule with that same severity that is specified for that policy triggers. For example, you can have rules specified for a policy that have a severity state of Critical. When any of the critical rules are triggered, all of the actions that are specified for the Critical severity state are executed.

Monitor task types

In addition to the standard Symantec Management Platform task types, Monitor-specific task types are available. Monitor-specific task types work the same as other tasks in Symantec Management Platform, but you may find that the Monitor task types are more useful.

See “[About Monitor Solution tasks and actions](#)” on page 117.

Predefined tasks are included in the Monitor packs. You can customize tasks further so that the execution of tasks is more useful.

You can add tasks to rules.

See “[Adding actions to rules](#)” on page 124.

You can also add tasks to policies.

See “[Adding actions to monitor policies](#)” on page 125.

Table 8-1 Monitor task types

Task type	Description
NT event	<p>Logs the events to the system that are then viewable in the Event Viewer.</p> <p>The task contains the following fields:</p> <ul style="list-style-type: none">■ Event Source – The message that displays in the Event Viewer.■ Event type – The severity level of the event, either informational, error, or warning.■ Category ID – You can enter a category ID of your choice or leave this field blank. The value you specify for category ID is displayed in the event information and can be sorted or searched on.■ Event ID – You can enter an event ID of your choice or leave this field blank. The value you specify for event ID is displayed in the event information and can be sorted or searched on.■ Parameter – The details of an NT event. When you enter text in this field it is always displayed with the events. If you want to display current Monitor data that is gathered at runtime, you can enter tokens in this field.
Syslog	<p>Logs the events to a UNIX system that are then viewable in the syslog file.</p> <p>The task contains the following fields:</p> <ul style="list-style-type: none">■ Indentation – The string that is prepended to every message, and is typically set to the program name.■ Priority – The priority level of the event, either emergency, alert, critical, error, warning, notice, info, or debug.■ Facility – An informational field that is associated with a syslog message. The syslog protocol defines it. It is meant to provide an indication from what part of a system a message has originated from.■ Message – The details of a syslog event. When you enter text in this field it is always displayed with the events. If you want to display current Monitor data that is gathered at runtime, you can enter tokens in this field.

Table 8-1 Monitor task types (*continued*)

Task type	Description
Process control	<p>Terminates a process or sets the priority of a process on a client computer. For example, you might want to stop notepad.exe from using too many system resources. You can specify the process as notepad.exe and set the priority for this task to Low. The task would check for running instances of notepad.exe and save system resources by setting that process to a lower priority. With notepad.exe at a lower priority, system resources are used for higher priority tasks.</p> <p>The task contains the following fields:</p> <ul style="list-style-type: none"> ■ Command – Lets you select Terminate to end the process, you can select Set Priority to adjust the priority level of the process. ■ Process Name – Lets you enter the name of the process you want to terminate or adjust the priority of. ■ Priority – If you selected Set Priority in the Command field, this drop-down list is enabled. Select the appropriate priority. ■ Apply command to all children – Lets you apply this command to all children of the process. For example, you can terminate the Internet Explorer process. If the computer had multiple tabs open in Internet Explorer, all tabs (instances) are terminated. If you do not select this check box, only the first tab is terminated. ■ Apply command to all instances – Lets you apply this command to all instances of the process. For example, you can terminate the Internet Explorer process. If the computer has multiple instances of Internet Explorer running, all of those instances are terminated. If you do not select this check box, only the first instance that is found is terminated.
Reset monitored resource	<p>Resets the Monitor Plug-in state.</p> <p>Resetting the Monitor Plug-in state affects Monitor behavior as follows:</p> <ul style="list-style-type: none"> ■ All of the rules that the plug-in knows about are reset to a normal severity state. ■ No rules are triggered. <p>You might want to run this task if the Notification Server computer and the Monitor Plug-in do not seem to be synchronized.</p>

Table 8-1 Monitor task types (*continued*)

Task type	Description
Poll metric on demand	Polls the specified metrics for a monitored resource or resources. When the task is run, the monitored resource or resources metrics are polled immediately. You can specify agent-based or agentless metrics for polling.

Adding tokens to a Send Email task

You can add Monitor tokens to a **Send Email** task. By adding tokens to a **Send Email** task, you make sure that you have the events that contain the information which is most useful to you. In the **Send Email** task you can add tokens to the body of the email or the subject field. Tokens are included every time the task is executed. You can have the same text appear when a **Send Email** task executes. By adding tokens to a **Send Email** task, you can gather monitor information at runtime and display it in the event. Monitor Solution specifies the available Monitor tokens that you can copy and place in your task definition.

See “[About Monitor Solution tasks and actions](#)” on page 117.

To add tokens to a Send Email task

- 1 Create or edit a rule.
- 2 Select a task from the **Task server** section.
Click the **New** symbol to create a server-side task.
- 3 In the **Create New Task** dialog box, at the bottom of the task listing page, click **Send Email**.
- 4 Click **OK**.
- 5 In the **Task Configuration** dialog box, click **Show tokens**.
A list of all available Monitor tokens is displayed.
See “[Monitor client and server token types](#)” on page 127.
- 7 In the **Task Configuration** dialog box, click **Edit task**.
- 8 Paste the tokens in the body of the email.
You can also add a token to the subject field to provide more clarity to the source of the alert.
- 9 (Optional) Repeat steps 1 through 8 for any additional tokens that you want to add.

10 Click **Save changes**.

After you close the task, in the **Task Configuration** dialog box, you can review a list of tokens that you chose.

11 Click **OK**.

Adding actions to rules

Rules are triggered when monitored metric data reaches a determined value or goes beyond an acceptable value range. When a rule is triggered an alert is raised. When an alert is raised the severity state of the monitored resource is changed to the severity setting of the rule that has triggered.

The severity state that you specify for a rule is reflected in the alert that is sent to the Event Console.

See “[About severity states](#)” on page 119.

Metric collect rules and metric rules interact with actions differently. Metric collect rules collect and forward data. Metric rules collect data and then evaluate it against the values you have specified in the rule. If the evaluation result is true, the rule is triggered and any actions specified for that rule are executed. Actions cannot be added to the metric collect rules because they do not support actions. With metric collect rules, there is nothing to evaluate, so a rule would never be triggered.

To add actions to rules

- 1 Create or edit a rule.

See “[Creating and editing rules in the rule library](#)” on page 108.

- 2 Select a task from one or both of the following options:

Task server section

Click the **Add** symbol to select a server-side task. Choose a task and click **OK**.

Task server tasks are run from the Task Server and can be run only if the Notification Server computer is reachable.

Monitor plug-in section

Click the **Add** symbol to select a task to run locally. Choose a task, and click **OK**.

Monitor Plug-in tasks are run from the Monitor Plug-in and can be run even if the Notification Server computer is not reachable. Monitor Plug-in tasks can only include client tasks.

- 3 If you want to include additional Monitor-specific information in any of the tasks, click the **Edit** symbol and configure the task.

See “[Adding tokens to a Send Email task](#)” on page 123.

- 4 (Optional) Repeat steps 1 through 3 for each task that you want to add.

- 5 The tasks are run in the order that they are displayed in the table. To change the task sequence, select a task and use the up and down arrows to place the tasks in the order that you require.

- 6 Click **OK**.

Adding actions to monitor policies

When monitored metric data reaches a determined value or goes beyond an acceptable value range, rules are triggered. A triggered rule sends an alert and changes the severity state. A policy’s task that is assigned a certain severity state is executed when a rule in the policy with a corresponding severity state is triggered.

See “[About Monitor Solution tasks and actions](#)” on page 117.

To add actions to monitor policies

- 1 Create or edit a monitor policy.

To create a monitor policy See “[Creating a monitor policy with the monitor policy wizard](#)” on page 63. to create a monitor policy.

To edit an agent-based monitor policy See “[Editing agent-based monitor policies](#)” on page 66. to edit an agent-based monitor policy

To edit an agentless monitor policy See “[Editing agentless monitor policies](#)” on page 67. to edit an agentless monitor policy

- 2 On the monitor policy page, on the **Actions** tab, select a severity state option from the top of the tab.

For each severity type, you can add an associated task or tasks that are executed when a resource changes to each severity type. For example, if a rule with a critical severity level is triggered, then all the tasks you specify for the critical severity level are executed.

See “[About severity states](#)” on page 119.

- 3 Select a task from one or both of the following options:

Task server section Click the **Add** symbol to select a server-side task. Choose a task and click **OK**.

Task server tasks are run from the Task Server, and can only be run if the Notification Server computer is reachable.

Monitor plug-in section Click the **Add** symbol to select a task to run locally. Choose a task and click **OK**.

Monitor Plug-in tasks are run from the Monitor Plug-in and can be run even if the Notification Server computer is not reachable. Monitor Plug-in tasks can only include client tasks.

This section is not available for agentless monitor policies.

Tasks can be run from the Task Server or tasks can be run locally if a Monitor Plug-in is installed on the monitored computer. Agentless monitor policies are only capable of running Task Server tasks.

- 4 Repeat the previous step to add all of the tasks that you require in the policy.

- 5 To include any additional monitoring-specific information in any of the tasks, click the **Edit** option and configure the task.
See “[Adding tokens to a Send Email task](#)” on page 123.
- 6 The tasks are run in the order that they appear in the table. To change the task sequence, select a task and use the up and down arrows to place the tasks in the order that you require.
- 7 (Optional) Repeat steps 1 through 6 for each severity state.
- 8 Click **Save changes**.

Monitor client and server token types

Tokens let you add additional Monitor information to actions. Specifying tokens for actions make the notifications that are sent to the Event Console more meaningful because they include Monitor-specific information in the event. The tokens are replaced with readable values after the task executes.

See “[About Monitor Solution tasks and actions](#)” on page 117.

Monitor Plug-in actions can use client tokens.

Table 8-2 Monitor Plug-in client tokens

Client tokens	Description
MONITOR_AGENT_NAME	The name of the computer with the Monitor Plug-in installed.
MONITOR_AGENT_STATE	The triggered alert on the plug-in that is the highest severity. The rule state and agent state are Monitor-specific and may not reflect the server's state in the Event Console.
MONITOR_ALERT_ID	The rule GUID or ID > Source > Category for Template rules.
MONITOR_CATEGORY_NAME	The name of the triggered rule's category.
MONITOR_EVENT_TIME	The time the rule triggered.
MONITOR_IN_MAINTENANCE_WINDOW	Reflects whether the plug-in is in a maintenance window. True if the plug-in is in a maintenance window, false otherwise.
MONITOR_INSTALL_DIR	The installation directory of the Monitor Plug-in.

Table 8-2 Monitor Plug-in client tokens (*continued*)

Client tokens	Description
MONITOR_METRIC_INFO	An XML fragment that describes the values that triggered the rule.
MONITOR_POLICY_NAME	The name of the triggered rule's policy.
MONITOR_OVERTIME_VALUE	The overtime setting for this rule.
MONITOR_POLICY_GUID	The GUID of the triggered rule's policy.
MONITOR_PREV_RULE_STATE	The previous rule state. The rule state and plug-in state are Monitor-specific and may not reflect the server's state in the Event Console.
MONITOR_RESOURCE_GUID	The resource that triggered the rule.
MONITOR_RULE_GUID	The GUID of the rule that triggered.
MONITOR_RULE_NAME	The name of the rule that triggered.
MONITOR_RULE_STATE	The severity state of the rule or "normal" if the rule is acknowledged. This rule state and Plug-in state is Monitor-specific and may not reflect the server's state in the Event Console.

Task server tasks can use all of the client tokens that are listed in the previous table, along with server tokens.

Table 8-3 Monitor Solution server tokens

Server tokens	Description
MONITOR_METRIC_INFO_HTML	Displays the XML metric information as HTML for email tasks.
MONITOR_SOURCE_GUID	The resource GUID of source computer.
MONITOR_TARGET_GUID	The resource GUID of target computer.
MONITOR_SOURCE_NAME	The name of the source computer.
MONITOR_TARGET_NAME	The name of the target computer.
MONITOR_SOURCE_DOMAIN	The domain name of the source computer.
MONITOR_TARGET_DOMAIN	The domain name of the target computer.

Table 8-3 Monitor Solution server tokens (*continued*)

Server tokens	Description
MONITOR_SOURCE_IP_ADDRESS	The IP address of the source computer.
MONITOR_TARGET_IP_ADDRESS	The IP address of the target computer.

Viewing Monitored Data

This chapter includes the following topics:

- [About the Monitoring and Alerting home page](#)
- [Viewing historical performance data](#)
- [Viewing real-time performance data](#)
- [Monitored Resources dialog box and Resources with Historical Data dialog box](#)
- [Viewing Monitor Solution reports](#)

About the Monitoring and Alerting home page

Monitor Solution lets you monitor the state of your entire enterprise in a single view through the **Monitoring and Alerting** home page. The home page makes it easy to check and ensure that all monitored computers and applications function properly.

The **Monitoring and Alerting** home page includes the following Web parts:

- **Launch Performance Viewer** – Used to enter the name of a computer and run the performance viewer.
See “[Viewing real-time performance data](#)” on page 133.
- **Monitored Resources by Status** – Shows a chart of monitored resources. The chart is organized according to severity status. The state of a computer is the most severe state of any triggered rule on the computer.
For example, if one rule state is warning and another is critical, the overall state of the computer is critical. If all rule states are normal, and then one rule state changes to warning, the computer state is set to warning. This Web part also shows computers with the Monitor Plug-in installed. You can select a

computer and launch the Performance Viewer, the Resource Manager, or the Event console.

- **Monitor Site Servers Status** – Shows a list of Monitor Site Servers and their Status.
- **Group View - Aggregate health by resource** – Shows the aggregate health of the devices and computers in your organizational groups.
- **Event Console** – Shows a consolidated view of all alerts that are raised.

Viewing historical performance data

The historical performance viewer is a console included with Monitor Solution that lets you view historical performance data. Historical data is available from both the Monitor Plug-in and the Remote Monitor Server.

To view historical performance data

- 1 In the Symantec Management Console, on the **Actions** menu, click **Monitor > Historical**.
- 2 Click the icon next to the **Device** field, and select a device that has historical data.
See “[Monitored Resources dialog box and Resources with Historical Data dialog box](#)” on page 134.

- 3 In **From** and **To**, specify the time period for which you want to view data.

The time period that you specified in **From** and **To** may contain no data in the beginning or at the end of the period. In this case **Summarized View** shows only the actual time when the data is available. The empty time line with no data in the beginning or at the end of the chart is not displayed.

- 4 To specify the metric data that you want to view, click **Metrics**, and use the **Available Metrics** dialog box, and then click **OK**.

- 5 In the **Summarized View**, drag the mouse across the graph to specify a range that you want to view.

- 6 In the **Detailed View**, select a point on the graph.

If available, the data that was last gathered for the selected point is displayed in the following sections: **Processes**, **Events**, **Ports**, and **Text Data**. The **Metrics** section continues to display the average, minimum, and maximum values for the whole range of data that is displayed in the **Detailed View**. However, the **Last Value** and **Last Time** columns in the **Metrics** section display the value at the selected point. If the selected point has no value, these columns display the value that precedes this point. If no value is available for the metric in the **Detailed View**, the **Last Value** and **Last Time** columns are left blank in the **Metrics** section.

See “[Viewing real-time performance data](#)” on page 133.

Viewing real-time performance data

The Performance Viewer is a console included with Monitor Solution that lets you view real-time performance data. Performance data is available from both the Monitor Plug-in and the Remote Monitor Server.

To view real-time performance data

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the **Launch Performance Viewer** Web part, either enter a name of the computer or use the computer browser to choose a computer.
- 3 In the **Registered Metrics** dialog box, select the metric data that you want to monitor, and then click **OK** (the limit is 50).

The performance viewer begins monitoring the computer and displays the following information:

- **Graph** – Displays graphical performance data. The data is scaled to fit within the limits of the graph. If you place the mouse pointer over a point on a graph line, the monitored metric data displays next to the mouse pointer. If you monitor multiple instance metrics, each instance has a separate graph line. You can use the **select metrics** option to monitor different metrics.
- **Metrics** – The metrics section displays all numeric metric data that is monitored.
- **Processes** – Displays the processes currently running on a monitored computer.

- **Events** – Displays All Windows NT event data.
- **Ports** – Displays the status of the monitored ports on the computer.
- **Text Data** – Displays the retrieved text data for command, custom DLL, custom COM object, WS-MAN, SNMP, SQL, and string-type Windows Management Instrumentation (WMI) metrics. The predefined WMI metrics are the only metrics that collect this type of data. If you create or use a custom DLL, COM object, SNMP, or command metric that retrieves this data, it is also displayed in this section.

See “[Viewing historical performance data](#)” on page 132.

Monitored Resources dialog box and Resources with Historical Data dialog box

The real-time and historical performance viewers use these dialog boxes to select a resource for which to view data.

See “[Viewing historical performance data](#)” on page 132.

See “[Viewing real-time performance data](#)” on page 133.

These dialog boxes provide the following data about the monitored resources:

- Device
- Domain
- IP address
- MAC address

Viewing Monitor Solution reports

Monitor Solution includes several reports that let you view data about your monitored computers.

See “[About the Monitoring and Alerting home page](#)” on page 131.

To view Monitor Solution reports

- 1 In the Symantec Management Console, on the **Home** menu, click **Monitoring and Alerting**.
- 2 In the left pane, click **Monitoring and Alerting > Monitor > Reports**.
- 3 Browse through the folders and click the report that you want to view.

Using alert management

This chapter includes the following topics:

- [About alerts](#)
- [About alert management](#)
- [Note on time zones and alerts](#)
- [About Event Console alert filters](#)
- [Alert Filter Settings page](#)
- [Filtering alerts](#)
- [Creating and saving alert filters](#)
- [About advanced search filters](#)
- [Creating advanced search filters](#)
- [Viewing alerts](#)
- [Hiding resolved alerts](#)
- [Alert Rule Settings page](#)
- [Creating an alert matching rule](#)
- [Adding or editing rules to discard alerts](#)
- [Forwarding alerts to another management system](#)
- [Running a task in response to an alert](#)
- [About Event Console tokens](#)
- [Event Console token types](#)

- [About the Event Console workflow rule](#)
- [About workflow rule configuration](#)
- [Adding or editing workflow rules](#)
- [About alert purging](#)
- [Purging old and low-severity alerts](#)
- [Viewing the health of an organizational group](#)
- [Working with Event Console tasks](#)
- [Change alert status task page](#)
- [Create resource task page](#)
- [Event Console purge policy task page](#)
- [Raise message task page](#)
- [Reprioritize alert task page](#)

About alerts

Alerts are the status messages that contain information about device or network health. Computers and many other devices can generate status messages using standard monitoring protocols, such as SNMP. The Symantec Management Platform collects and tracks these status messages.

Each status message that is received is converted from its native format into a common format that is called an alert. During conversion, alerts are associated with the affected resource in the CMDB and are assigned a severity and a status. Severity ranges from normal to critical, and alert status can be new, acknowledged, or resolved.

Alerts from multiple protocols are displayed using common severity and status. All received alerts are displayed in the Event Console.

See “[Viewing alerts](#)” on page 145.

See “[About Event Console alert filters](#)” on page 138.

About alert management

Alert management shows a consolidated view of device health across your network. You can view health by network layout, organizational group, or by directly monitoring the list of received alerts in the Event Console.

The Event Console reduces the need to maintain separate tools to monitor computers, software, printers, and other devices. The Event Console collects SNMP traps and other status messages and displays them in a single location. All status messages are converted to a common format that links each received message to the affected resource in the CMDB. These formatted messages are called alerts.

See “[About alerts](#)” on page 136.

Advanced search features let you find specific alerts or groups of alerts quickly, using search operators similar to those found in common applications.

See “[About advanced search filters](#)” on page 142.

See “[Creating advanced search filters](#)” on page 144.

The Event Console also provides a rule-based triggering system that lets you process alerts in the following ways by creating alert matching rules:

- Discard specific alerts from the database.
See “[Adding or editing rules to discard alerts](#)” on page 148.
- Forward alerts to another management system.
See “[Forwarding alerts to another management system](#)” on page 149.
- Execute task server tasks in response to specific alerts.
See “[Running a task in response to an alert](#)” on page 149.
- Initiate a workflow in response to specific alerts.
See “[Adding or editing workflow rules](#)” on page 153.

See “[Viewing alerts](#)” on page 145.

See “[Creating an alert matching rule](#)” on page 147.

Note on time zones and alerts

Note: If the Notification Server computers and the SQL Server computers are not set to the same time and the same time zone, then any alerts that have occurred in the past few hours are not displayed in the Event Console. See “[Creating and saving alert filters](#)” on page 141.

About Event Console alert filters

The Event Console in Symantec Management Platform displays alerts in a grid layout. This grid may contain thousands of alerts. Alert filters let you sort the alerts so that you can analyze and manage them. You access this grid from Symantec Management Console when you click **Manage > Events and Alerts**.

The Event Console in Symantec Management Platform contains several rule types that represent automated, event-based actions. The rule types include discarding, forwarding, task, and workflow rules. Discarding rules filter and discard matching alerts. Forwarding rules forward a Simple Network Management Protocol (SNMP) trap to a downstream listener. Task rules initiate Symantec Management Platform task server tasks. Before version 7.1 of the platform, a direct way to initiate a deployed workflow process was unavailable. With the addition of a workflow rule in version 7.1 of the Event Console, an event can automatically start a workflow process. This workflow process can pass along valuable event data.

Previous versions of the platform let you filter alerts into manageable subsets. However, before version 7.1 you did not have the option to save and re-use those filters. Beginning with version 7.1, you can create, save, and re-use filters.

See “[About alerts](#)” on page 136.

A new function with version 7.1 now lets you use advanced filters to manage alerts. The advanced filter function is available from the Event Console grid.

See “[About advanced search filters](#)” on page 142.

The following alert filtering tools are available on the main **Event Console** page:

- A drop-down list of predefined filters. You can click **Actions** to see a list of available filtering actions. You can also search by one of the following alert criteria:
 - **Alerts in last 24 hours**
 - **Alerts in last 7 days**
 - **Critical Alerts in last 24 hours**
 - **Critical Alerts Only**
 - **Exclude Informational Alerts**
 - **Exclude Monitor Alerts**
 - **Informational Alerts Only**
 - **Major Alerts Only**
 - **Monitor Alerts Only**

- **Normal Alerts Only**
- **Undetermined Alerts Only**
- **Warning Alerts Only**
- A color-coded, left-click **Status Progress Bar** control. This control lets you see the number of alerts by severity level, as follows:

Violet	Undetermined
Yellow	Warning
Orange	Major
Green	Normal
Blue	Informational
Red	Critical

You can access the color-coded status progress bar control using a left-mouse click. This bar appears in the **Alerts** pane. When you click a color section on the status bar, the grid view changes. The view shows only those alerts that match the severity level of the color that you clicked. For example, if you click yellow on the status bar, then the grid shows alerts with severity **Warning**. After you filter by severity level, you may have to select **Exclude Informational Alerts** or **Monitor Alerts Only** to see the complete list of alerts again.

See “[Filtering alerts](#)” on page 140.

- A status bar that presents the following icons:
 - **Details.** Opens the **Alert Details** dialog box for the selected alert.
 - **Acknowledge.** Lets you acknowledge a selected alert. In the **State** column, a blue flag indicates an acknowledged alert.
 - **Resolve.** Flags the selected alert with a check mark in the **State** column. When you right-click a resolved alert, you can view alert details. You can also view the available rules for discarding the alert or open the **Resource Manager** in a new window.
If you click **Discarding Rules** with a resolved alert selected, you can create a global discard filter rule or create a resource discard filter rule.
 - **Actions.** When you select an alert and click the down-arrow next to this icon, you see the options that were listed previously in this list. You also see one addition.
When you click any alert, you can manage it by changing its severity to any of the following:

- **Major**
- **Warning**
- **Informational**
- **Undetermined**
- **Normal**
- **Critical**
- An **Alert Filter Settings** page for managing the filters that you save and reuse.
A tool icon next to the predefined filters drop-down list opens the **Alert Filter Settings** page. This page is where you can create filters and save them for re-use.
See “[Alert Filter Settings page](#)” on page 140.
See “[Creating and saving alert filters](#)” on page 141.
- A search field that lets you enter custom search criteria.
The magnifying glass icon next to the search field opens the **Advanced Search** pane.
See “[About advanced search filters](#)” on page 142.

Alert Filter Settings page

The **Event Console Alert Filter Settings** page lets you manage the filters that can then be applied to the alert grid. You can add new filters with specific alert filter conditions, edit existing filters, or delete filters.

See “[About Event Console alert filters](#)” on page 138.

In the Symantec Management Console, you access this page from the **Settings > All Settings** menu. After you expand **Settings > Monitoring and Alerting**, you can click the page title, and the page opens on the right.

See “[About alerts](#)” on page 136.

When you add, edit, and delete alert filters, you may also need to work with alert rules.

See “[Alert Rule Settings page](#)” on page 146.

See “[Creating an alert matching rule](#)” on page 147.

Filtering alerts

The Event Console grid can contain thousands of alerts, which you can filter. If the alerts that you expect to see are not displayed, they may be hidden, or a filtering

rule has blocked them. For example, some administrators prefer to hide warning alerts.

See “[About Event Console alert filters](#)” on page 138.

See “[Creating and saving alert filters](#)” on page 141.

See “[About advanced search filters](#)” on page 142.

In the Event Console, the default filter is **Exclude information alerts**. When you open the alert grid, this default filter is applied. Anytime you click the **Refresh** icon in the browser window, the selected filter is reset. You can also clear filters and select new ones.

To filter alerts

- 1 In the Symantec Management Console, on the **Manage** menu, click **Events and Alerts**.

If the filter you see in the filter drop-down box is not the one you want to use, perform the next step.

- 2 (Optional) In Event Console click the X icon to the right of the filter drop-down box, or delete the filter text from the box.

The alerts are cleared from the grid, and **Select a filter** appears in the drop-down box.

- 3 Click the down-arrow next to the drop-down box to select a different filter.

As soon as you select a different filter from the drop-down list, the grid view changes. It shows only the alerts that pertain to the selected filter. You can click any other control on the page, except **Refresh**, and the filter that you chose remains active. If you need to view alerts for more than one filter, you can open multiple instances of Event Console. You then select a different filter in each window.

Creating and saving alert filters

Before Symantec Management Platform version 7.1, you filtered alerts using the advanced filter functionality that was built in to the Event Console. However, you did not have the option of saving your filters. Beginning with Symantec Management Platform 7.1, you can create, save, and re-use filters.

See “[About Event Console alert filters](#)” on page 138.

See “[Filtering alerts](#)” on page 140.

To create and save alert filters

- 1 In the Symantec Management Console, on the **Manage** menu, **Events and Alerts**.
- 2 In the **Event Console** window, click the **Tools** icon to the right of the filters drop-down list.
- 3 In the **Alert Filter Settings** dialog box, click **Add** to create a new filter.
- 4 Click the default filter name to give the filter a unique, descriptive name.
- 5 Under the filter name, click **New filter description**, and enter a description of what the alert can filter.
- 6 Under **Filter Condition**, under **This filter evaluates the following conditions**, click each drop-down list to set the conditions that the new filter should evaluate.
Click **Add** to add multiple conditions for a single filter to evaluate.
- 7 In the right corner of the **Status** pane, click **On** to enable this alert for use.
If you want to create the filter but not enable it, leave the status set to **Off**.
- 8 Click **Save**.
After you close the **Alert Filter Settings** window, the Event Console drop-down list includes the new filter.

About advanced search filters

The ability to perform advanced searches using alert filters is new with Symantec Management Platform 7.1. You can use this built-in search function to help you manage alerts. The advanced filtering function is available from the main **Event Console** window, which you access from Symantec Management Console under **Manage > Events and Alerts**.

See “[About alert management](#)” on page 137.

In the main **Event Console** window, you click the magnifying glass next to the **Search** field and the **Advanced Search** pane opens. In the **Advanced Search** pane, you can choose from a predefined list of search criteria or type your own criteria. You can add other rules to an advanced search to further customize it.

See “[Filtering alerts](#)” on page 140.

See “[Creating advanced search filters](#)” on page 144.

The following advanced search tools let you narrow the list of filters to manage:

- Three drop-down lists from which you select subsets of alerts

In the first drop-down list you can enter or select a search type.	Search types include the following: <ul style="list-style-type: none">■ Category■ Count■ Description■ First occurred■ Host■ Last occurred■ Severity■ State
In the second drop-down list you can enter or select a search operator.	The search type that you select from the first drop-down list determines the search operators that appear in the second drop-down list. Some or all of the following search operators appear in the second drop-down list: <ul style="list-style-type: none">■ Equals■ Not equals■ Contains■ Less than■ Greater than
In the third drop-down list, enter or select additional search criteria to apply to the selected search type that uses the selected search operator.	After you select or enter a search type and one or more search operators, additional search criteria appear in the third drop-down list. For example, if you enter or select count > greater than in the first two drop-down lists, you can select a value. You can select a value such as 5 to view only those alerts that have occurred more than five times. Or, if you selected Host > equals from the first two drop-down lists, then you can select from a list of computers.
<ul style="list-style-type: none">■ An Add Rule option that lets you access a drop-down list and add the following operators to your search:<ul style="list-style-type: none">■ AND■ OR■ NOT	

- A color-coded, left-click status progress bar above the **Advanced Search** pane. This control lets you filter alerts by severity level. After you filter by severity level, you may have to select **Exclude Informational Alerts** or **Monitor Alerts Only** to see the complete list of alerts again.
- See “[About Event Console alert filters](#)” on page 138.

An advanced search lets you view the same types of information that you can view about all alerts:

See “[Viewing alerts](#)” on page 145.

When you click any alert, you can manage it by changing the state. Click any flag in the **State** column to view details about an alert, acknowledge the alert, resolve it, or perform another action. These actions are accessible from the **Actions** drop-down list above the **Advanced Search** pane.

Creating advanced search filters

The advanced search feature in the Event Console lets you find specific alerts or groups of alerts quickly. The search operators that you can use are similar to the operators that are found in common applications. One such feature is advanced search, which lets you create filters to narrow your search.

See “[About advanced search filters](#)” on page 142.

To create advanced search filters

- 1 In the Symantec Management Console, on the **Manage** menu, click **Events and Alerts**.
- 2 On the status bar, click the **Advanced Search** symbol.
- 3 In the **Advanced Search** pane, to the right of **Where:**, click the down-arrow in each drop-down search list. Select the criteria by which you want to search.
- 4 To focus your search further, click the **Add Rule** down-arrow.
- 5 Select the operator that you want to use to focus your search.

The following options are available:

- **AND**
- **OR**
- **NOT**

- 6 (Optional) In the new row of drop-down search lists, click the down-arrows to select additional criteria.

- 7 (Optional) Continue to add search criteria until you are satisfied that the search returns the results that you want.
- 8 Click **Search**.
- 9 (Optional) To clear part or all of the search criteria, click **Clear**.
- 10 (Optional) To close the **Advanced Search** pane, click **Close**.
If you close the pane, your search criteria are cleared. To collapse the pane without losing the search criteria, click the up-arrow next to **Close**.

Viewing alerts

Alerts are collected and displayed in the Event Console. If expected alerts are not displayed, they might be hidden or a filtering rule might be configured to block them.

See “[About alerts](#)” on page 136.

See “[Hiding resolved alerts](#)” on page 146.

You can view the following information about each alert:

- **Severity**
- **State**
- **Host**
- **First occurred**
- **Last occurred**
- **Count**
- **Category**
- **Definition**
- **Protocol**
- **Description**

Status flags in the **State** column let you quickly see additional alert information.

- **Details**
- **Acknowledge**
- **Resolve**
- **Actions**

To view alerts

- ◆ In the Symantec Management Console, on the **Manage** menu, click **Events and Alerts**.

Hiding resolved alerts

Hiding alerts prevents them from being displayed on the Event Console alert grid. Hiding alerts lets you remove the alerts that have been resolved. Hiding resolved alerts lets you focus on unresolved alerts. Hidden alerts remain in the alert database until they are removed by purging.

See “[About alerts](#)” on page 136.

See “[Purging old and low-severity alerts](#)” on page 155.

To hide resolved alerts

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Event Console Settings**.
- 3 On the **Event Console Settings** page, specify a time for resolved alerts to remain viewable in Event Console. After this interval, the resolved alerts are hidden automatically.
- 4 Click **Save changes**.

Alert Rule Settings page

The **Event Console Alert Rule Settings** page lets you add different rules for several purposes. You can create some rules that discard or forward alerts. You can also create some task rules and rules for initiating workflow tasks.

In the Symantec Management Console, you access this page from the **Settings** menu click **All Settings**. After you expand **Settings > Monitoring and Alerting** and then you click **Alert Rule Settings**.

The available tabs on this page are as follows:

■ Discarding Rules

See “[Adding or editing rules to discard alerts](#)” on page 148.

■ Forwarding Rules

See “[Creating an alert matching rule](#)” on page 147.

■ Task Rules

See “[Creating an alert matching rule](#)” on page 147.

■ **Workflow Rules**

See “[Adding or editing workflow rules](#)” on page 153.

Creating an alert matching rule

Alert matching rules contain conditions, such as alert type or date received, to identify specific alerts as the Event Console receives them. These rules are used when you discard or forward alerts, execute tasks, or initiate a workflow.

See “[Adding or editing rules to discard alerts](#)” on page 148.

See “[Running a task in response to an alert](#)” on page 149.

See “[Forwarding alerts to another management system](#)” on page 149.

See “[Adding or editing workflow rules](#)” on page 153.

You can match alerts by type, severity, affected resource, and many other criteria.

To create an alert matching rule

- 1 In Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Alert Rule Settings**.
- 3 On the **Alert Rule Settings** page, click the tab that corresponds to the type of rule you want to create.

The following rule types are available:

- **Discarding Rules**
- **Forwarding Rules**
- **Task Rules**
- **Workflow Rules**

- 4 In the left pane, click **Add** to create a new alert matching rule.
- 5 In the alert rule builder on the right, click the generic rule name and create a unique name for the rule.
- 6 Under the rule name, click the generic description and update the text to describe the new rule.
- 7 Click **Add** to define the rule conditions.

8 Define matching criteria for the conditions.

You can re-order conditions and move them up and down or left and right to create nested evaluations. During evaluation, nested evaluations are performed first.

- 9** If you create a new workflow rule, define the workflow to run when a matching alert is received.
- 10** At the upper right of the page on the **Status** bar, click the colored circle, and then click **On** to enable the rule.

The default status is **Off**.

- 11** Click **Save**, and then check the rule builder for any error or any warning messages.

Adding or editing rules to discard alerts

You may need to delete incoming alerts under certain conditions. Or, you may want to delete duplicate alerts. In the **Event Console Alert Rule Settings** page, you can create an alert matching rule to discard the alerts that meet your criteria. These alerts are removed as soon as they are received and are not imported into the Configuration Management Database.

To optimize performance of the platform, and Notification Server in particular, you should create discard rules to remove redundant alerts. The **Discarding Rules** tab lets you configure multiple conditions for the incoming alerts that the system should discard.

See “[About alerts](#)” on page 136.

See “[Creating an alert matching rule](#)” on page 147.

Filtered alerts are never stored in the alert database and are unavailable when reports are generated. If you want to store alerts but do not want to display them in the Event Console, hide them instead.

See “[Hiding resolved alerts](#)” on page 146.

To add or edit a rule to discard an alert

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Alert Rule Settings**.
- 3 On the **Alert Rule Settings** page, click the **Discarding Rules** tab, and then click **Add** to create a new alert matching rule.

- 4 Define the matching conditions and the workflow to run when a matching alert is received.
- 5 At the upper right of the page, click the colored circle, and then click **On** to enable the rule.
- 6 Click **Save**.

Forwarding alerts to another management system

Alerts can be forwarded as SNMP traps to other management systems.

To forward alerts to another management system

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Alert Rule Settings**.
- 3 Click the **Forwarding Rules** tab, and then click **Add** to create a new alert matching rule.
See “[Creating an alert matching rule](#)” on page 147.
- 4 Define the matching conditions, and add the IP address or host name of the management system where the alerts that match the rule should be forwarded.
- 5 At the upper right of the page, click the colored circle, and then click **On**.
- 6 Click **Save**.

Running a task in response to an alert

Event console can perform a task server task in response to a received alert.

A single alert can trigger multiple, independent tasks. If multiple tasks must be performed in order, you can combine these tasks into a job.

See “[Creating an alert matching rule](#)” on page 147.

Running a task in response to an alert

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Alert Rule Settings**.
- 3 Click the **Task Rules** tab, and then click **Add** to create a new alert matching rule.

- 4 Define the matching conditions and the task to execute when a matching alert is received.

Select a task from both or one of the following options:

Create a new task.

Click the **New** symbol to create a new task.

You can add Event Console-based tokens to your new task.

See “[Event Console token types](#)” on page 150.

Use an existing task.

Click the **Add Existing** symbol to use one of the created existing tasks.

- 5 At the upper right of the page, click the colored circle, and then click **On**.
6 Click **Save**.

About Event Console tokens

Event Console tokens provide information when a task executes in response to a received alert. When a task executes, the Event Console tokens are replaced with readable values.

Note: Every solution has specific tokens that you can use to create a task for that solution. However, Event Console only resolves tokens to readable values in response to alerts from the tasks that were created using Event Console tokens.

For example, if you have a Monitor task that was created using Monitor tokens, you can assign it to Event Console. However, when Event Console receives an alert from that task, it cannot translate the tokens or generate readable values. Instead, the Monitor token returns its literal value rather than a readable value.

See “[Running a task in response to an alert](#)” on page 149.

See “[Event Console token types](#)” on page 150.

Event Console token types

Event Console tokens provide information when a task executes in response to a received alert. When a task executes, the Event Console tokens are replaced with readable values.

See “[About Event Console tokens](#)” on page 150.

Table 10-1 Describes the readable values that are associated with the Event Console tokens

Event Console tokens	Description
%!ALERTCATEGORYGUID!%	The GUID of the alert category.
%!ALERTDEFINITIONGUID!%	The GUID of the alert definition.
%!ALERTGUID!%	The GUID of the alert.
%!ALERTHOSTNAME!%	The host name or IP address of the resource that raised the alert.
%!ALERTMESSAGE!%	The message text of the alert.
%!ALERTPRODUCTGUID!%	The GUID of the product (in the case of a solution) that raised the alert.
%!ALERTPROTOCOLGUID!%	The GUID of the protocol that raised the alert.
%!ALERTRESOURCEGUID!%	The GUID of the NS Resource that raised the alert.
%!ALERTSEVERITYLEVEL!%	The enumeration value which represents the severity of the alert. <ul style="list-style-type: none">■ Critical = 50■ Major = 40■ Warning = 30■ Informational = 20■ Undetermined = 10■ Normal = 0
%!ALERTTIMESTAMP!%	The date and time the alert was raised.
%!ALERTVARIABLE!%	The variable name. Each variable from the alert is passed where "variable_name" is the name of the variable and the value is the variable value string.

About the Event Console workflow rule

In Symantec Management Platform, incidents are reported and presented in the **Event Console** window. Incidents are reported as alerts based on the filtering rules that you set up.

See “[Creating an alert matching rule](#)” on page 147.

Workflow rules let you forward received alerts into a deployed workflow. All information about alerts and their variables are passed into the workflow.

All enabled Event Console rules are evaluated against each inbound alert. Alerts that match all the conditions of a rule trigger that rule. If a rule is triggered, a Symantec Management Platform task can be initiated. Options on the **Workflow Rules** tab on the **Alert Rule Settings** page let you initiate workflow processes automatically when alerts are received in the Event Console. The **Workflow Rules** tab on the **Alert Rule Settings** page in the Event Console lists existing workflow rules and lets you add, edit, and delete rules.

See “[Adding or editing workflow rules](#)” on page 153.

The workflow rule includes a workflow rule configuration page.

See “[About workflow rule configuration](#)” on page 152.

The workflow rule includes the following alert severities that support the additional levels:

- Normal
- Undetermined
- Informational
- Warning
- Major
- Critical

When you click the **Workflow Rules** tab, you see the workflow rules that are listed in the left pane. When you click a rule, its description and details appear in the right pane. In this pane you can configure multiple conditions for incoming alerts and select a destination workflow. For more information about using workflows, see the *Workflow Solution User Guide*.

About workflow rule configuration

You configure rules in Event Console through tabs on the **Event Console Alert Rule Settings** page.

See “[Alert Rule Settings page](#)” on page 146.

The Workflow Rules tab lets you add, edit, enable, and delete workflow rules. Only the rules that are marked as On (enabled) are actively processed when Event Console receives new alerts.

See “[Running a task in response to an alert](#)” on page 149.

See “[About the Event Console workflow rule](#)” on page 151.

Event Console makes use of the Workflow Directory API. This API provides a complete list of available, deployed workflows. The API also specifies which workflow entry points are designed to be launched without a form interface; that is, by process. The entry point details include a Category field. This field lets Event Console determine if a Workflow rule calls the entry point, by design. Once complete, only the rules that are in the Process Start, Event Console category appear in the workflow selection drop-down list.

The workflow rule is secured with the following Symantec Management Platform security permissions:

- Enable Workflow Rules. This permission lets you enable and disable existing Workflow rule items in the **Alert Rule Settings** page.
- Modify Workflow Rules. This permission allows users to create, edit, and delete workflow rule items in the **Alert Rule Settings** page.

See “[Adding or editing workflow rules](#)” on page 153.

Adding or editing workflow rules

You can add, edit, or delete workflow rules. You add or edit workflow rules to forward all information about received alerts and their variables into a deployed workflow.

The **Workflow Rules** tab on the **Alert Rule Settings** page in the Event Console lists existing workflow rules. You can use the existing rules, or you can add, edit, and delete rules.

See “[About the Event Console workflow rule](#)” on page 151.

The Event Console **Workflow Rules** tab offers 14 rule condition types to filter the events that trigger the rule. These conditions are the same across all Event Console rule types (including discarding, forwarding, task, and workflow). No new functionality is provided to the conditions.

You can add or edit a rule to evaluate the following conditions:

- **Category.** The event category
- **Count.** The number of deduplicated alerts received (within a period of time)
- **Date.** The date on which the event occurred
- **Day of week.** The day of the week on which the event occurred
- **Definition.** The specific event type name
- **Host name.** The name or IP address of the resource
- **Message.** The event description text

- **Product.** The event-reporting product source
- **Protocol.** The protocol that is used to report the event
- **Resource.** The managed or unmanaged resource
- **Resource target.** The resource belonging to a specified group
- **Severity.** The severity level of the event
- **Time of day.** The time of day at which the event occurred
- **Variable.** All name data pairs or value data pairs that are provided in the event details

Adding or editing workflow rules

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Alert Rule Settings**.
See “[Alert Rule Settings page](#)” on page 146.
- 3 On the **Alert Rule Settings** page, click the **Workflow Rules** tab, and then click **Add** to create a new alert matching rule.
See “[Creating an alert matching rule](#)” on page 147.
- 4 In the rule builder on the right, click **Add** to create a new rule or click a rule that you want to edit.
- 5 Define the matching conditions for a new rule, or edit the conditions for an existing rule. Define the workflow to run when a matching alert is received.
- 6 At the upper right of the page, click the colored circle, and then click **On** to enable the rule.
- 7 Click **Save**.

About alert purging

Alert purging is a feature of Event Console that removes alerts from the database. Removal from the database is age-based. Age-based alert purging means that removal is based on the age of the alerts. Alert purging also lets you remove a target number of stored alerts and offers the enhanced function of purging unresolved alerts.

Age-based purging removes all alerts that are older than the specified number of days, which is calculated in 24-hour periods from the current time. Age-based purging removes old alerts regardless of their status or severity.

Target-number purging decreases the number of stored alerts by prioritizing the alerts that are based on age, status, and severity. When a target-number purge occurs, all resolved alerts that are older than the purge age are deleted first, from least to most severe. As soon as the number of stored alerts is less than the threshold, purging stops.

The "Do not purge unresolved alerts" function is enabled by default. If the threshold has not been reached when the purging is complete and you have disabled the "Do not purge unresolved alerts" function, unresolved alerts begin auto-resolving. Auto-resolved alerts are purged. If you have enabled the "Do not purge unresolved alerts" function, then purging is completed even if the threshold has not been met.

This purging process continues on the alerts that are newer than the specified purge age. Purging continues as long as needed to bring the number of alerts to less than the threshold. The system purges alerts by severity. Purging occurs in groups, not individually.

You can remove all the alerts from your database by first disabling the "Do not purge unresolved alerts" function. Then, you either set the target number of alerts to purge to zero (0) or set the age to zero days old for purging.

See ["Purging old and low-severity alerts"](#) on page 155.

Purging old and low-severity alerts

You should purge alerts periodically to maintain database performance.

See ["About alert purging"](#) on page 154.

To purge old and low-severity alerts

- 1 In the Symantec Management Console, on the **Settings** menu, click **All Settings**.
- 2 In the left pane, click **Settings > Monitoring and Alerting > Event Console Purging Maintenance**.
- 3 On the **Event Console Purging Maintenance** page, select the purge method that you want to use, and schedule the purge cycle.
- 4 Click **Save changes**.

Viewing the health of an organizational group

The group view in Symantec Management Console shows the aggregate health of the devices and computers in your organizational groups. If a device is not managed, alerts are not included when group health is displayed.

The group view is installed as part of Server Management Suite.

See “[About alerts](#)” on page 136.

See “[Viewing alerts](#)” on page 145.

To view the health of an organizational group

- 1 In Symantec Management Console, on the **Home** menu, click **Server Management Suite Portal**.
- 2 To view the group view, navigate to the Web part that is labeled **Group View - Aggregate health by resource**.

Working with Event Console tasks

You can create, modify, and delete Event Console tasks from a single location.

To work with Event Console tasks

- 1 In Symantec Management Console, on the **Manage** menu, click **Jobs and Tasks**.
- 2 In the left pane, click **System Jobs and Tasks > Monitoring and Alerting**, right-click **Event Console Tasks**, and then click **New > Task**.
- 3 In the **Create New Task** dialog box, in the left pane, click **Monitoring and Alerting > Event Console**.

In this location, you can work with the following Event Console tasks:

- **Change alert status task**
See “[Change alert status task page](#)” on page 157.
- **Create resource task**
See “[Create resource task page](#)” on page 157.
- **Event Console purge policy task**
See “[Event Console purge policy task page](#)” on page 157.
- **Raise message task**
See “[Raise message task page](#)” on page 157.
- **Reprioritize alert task**
See “[Reprioritize alert task page](#)” on page 158.

Change alert status task page

This page lets you acknowledge or resolve an alert that is received in Event Console. This task must be used with an alert rule that you set up to trigger when the alert is received.

See “[Working with Event Console tasks](#)” on page 156.

Table 10-2 Option on the Change alert status task page

Option	Description
New alert status	The status to assign to the matching alert.

Create resource task page

This page lets you create a new resource for a device that raises an alert but is unknown to the Notification Server computer . This task is useful as an Alert Rule action with an "Alert Resource is not managed" condition. A basic resource is created (a resource GUID is assigned to the host name of the alert that triggers the task).

See “[Working with Event Console tasks](#)” on page 156.

Event Console purge policy task page

This page lets you purge Event Console settings. However, all Event Console purge settings should be done through the purge settings user interface page in Event Console. This task is not meant for consumption through the Task Server user interface.

See “[Working with Event Console tasks](#)” on page 156.

Raise message task page

This page lets you raise a message within Notification Server. This task can be used with a task rule to cause an AlertRaisedNSMessage in Notification Server. However, this message occurs in Event Console when an alert is received without a task dependency. This task is not meant for consumption through the Task Server user interface.

See “[Working with Event Console tasks](#)” on page 156.

Reprioritize alert task page

This page lets you change the severity level of an alert. Use this task with an alert rule that you set up to trigger when the alert is received.

See “[Working with Event Console tasks](#)” on page 156.

Table 10-3 Option on the Change alert status task page

Option	Description
New alert severity	The severity to assign to the matching alert.

Index

A

advanced search filters
 about 142
 creating 144
agent-based monitor policies
 editing 66
agent-based monitoring
 preparing computers 33
agentless monitor policies
 editing 67
agentless monitoring 14
 about 47
agentless-based monitoring
 network discovery 48
Alert Filter Settings page 140
alert filters
 creating 141
 saving 141
alert management
 about 137
alert matching rule
 creating 147
Alert Rule Settings page 146
alerts
 about 136
 filtering 140
 forwarding 149
 forwarding to another management system 149
 purging 155
 running a task 149
 viewing 145
application detection 72
 about 68
application detection types 70

C

Change alert status task page 157
client tokens
 types 127
context-sensitive help 19
Create resource task page 157

D

database maintenance
 about 23
documentation 19

E

Event Console
 token types 150
Event Console alert filters
 about 138
Event Console purge policy task page 157
Event Console tasks
 working 156
Event Console tokens
 about 150
 types 150
Event Console workflow rule
 about 151

H

heartbeat
 about 26
heartbeat settings
 setting up 27
help
 context-sensitive 19
historical performance data
 viewing 132
historical performance viewer 14

I

import policy
 creating 22

M

metric evaluation
 about 111
metrics
 about 75
 creating 75

- metrics (*continued*)
 - editing 75
 - types 77
 - monitor actions
 - about 117
 - Monitor Pack for Servers
 - about 17
 - monitor packs 17
 - monitor policies 17
 - monitor packs 14, 17
 - about 22
 - importing 22
 - Monitor Plug-in 14
 - about 30
 - about configuring 31
 - about installing 30
 - configuration policies 31
 - configuring settings 35
 - creating settings 34
 - installation policies 30
 - installing 42
 - policies 30
 - profiling 32
 - uninstalling 44
 - upgrading 43
 - Monitor Plug-in configuration settings
 - Data Collection tab 39
 - General tab 36
 - Maintenance Windows tab 42
 - Performance Tuning tab 38
 - Monitor Plug-in settings
 - configuring 35
 - creating 34
 - monitor policies 14
 - about 31
 - adding actions 125
 - adding application detection 69
 - adding computers 72
 - adding rules 68
 - creating 65
 - creating with the wizard 63
 - monitor server
 - configuring 21
 - preparing 21
 - monitor server's heartbeat settings
 - setting up 27
 - monitor service
 - about 48
 - adding to a site server 53
 - monitor service (*continued*)
 - removing from a site server 52
 - monitor site server
 - configuring settings 55
 - data collection settings 59
 - general settings 55
 - performance settings 57
 - reports 54
 - Monitor Solution
 - about 13
 - components 14
 - features 16
 - metrics 16
 - new features 14
 - rules 16
 - tasks 16
 - what's new 14
 - Monitor Solution reports
 - viewing 134
 - monitor tasks
 - about 117
 - types 120
 - Monitored Resource dialog box 134
 - Monitoring and Alerting
 - home page 131
 - Monitoring and Alerting home page
 - about 131
 - multiple instance metrics
 - about 104
- N**
- New COM metric page 80
 - New Command Metric page 80
 - New Compound Metric page 84
 - New Custom DLL Metric page 85
 - New Group Metric page 86
 - New HTTP Metric page 87
 - New Log Event Metric page 89
 - New Metric Evaluation page 112–114
 - New Performance Counter Metric page 92
 - New Ping Metric page 91
 - New Port Metric page 93
 - New Smart Metric page 96
 - New SNMP Metric page 94
 - New SQL Metric page 96
 - New Windows Process Metric page 98
 - New Windows Service Metric page 99
 - New WMI Metric page 100
 - New WS-MAN Metric page 102

O

organizational group
viewing health 155

P

performance data
maintaining 25
Pluggable Protocols Architecture
installing 51
PPA. *See* Pluggable Protocols Architecture
purging alerts 155
about 154

R

Raise message task page 157
real-time performance data
viewing 133
real-time performance viewer 14
Release Notes 19
Remote Monitoring Server
configuring 55
data collection settings 59
general settings 55
performance settings 57
remote monitoring site server
setting up 49
Reprioritize alert task page 158
resolved alerts
hiding 146
Resources with Historical Data dialog box 134
rules
about 107
adding actions 124
adding metric 77
cloning 108, 110
creating 108, 110
editing 108, 110
types 115
rules to discard alerts
adding 148
editing 148

S

Send Email task
adding tokens 123
server tokens
types 127

severity states
about 119

W

workflow rule configuration
about 152
workflow rules
about 151
adding 153
configuration 152
editing 153

Altiris™ Monitor Solution for Servers 7.1 SP2 Symantec™ Third-Party Legal Notices

This appendix includes the following topics:

- [Third-Party Legal Attributions](#)
- [Expat XML Parser v2.0.1](#)
- [Net-SNMP v 5.4.1](#)
- [RegExp](#)

Third-Party Legal Attributions

This Symantec product may contain third party software for which Symantec is required to provide attribution (“Third Party Programs”). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. This appendix contains proprietary notices for the Third Party Programs and the licenses for the Third Party Programs, where applicable.

Expat XML Parser v2.0.1

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003 Copyright 1989, 1991, 1992 by Carnegie Mellon University Derivative Work - 1996, 1998-2000 Copyright 1996, 1998-2000 The Regents of the University of California. All Rights Reserved Copyright (c) 2001-2003, Networks Associates Technology, Inc. All rights reserved.

MIT License

This code is licensed under the license terms below, granted by the copyright holder listed above. The term "copyright holder" in the license below means the copyright holder listed above.

Copyright (c) <year> <copyright holders>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Net-SNMP v 5.4.1

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Copyright (c) 2001-2003, Networks Associates Technology, Inc

All rights reserved.

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.

All rights reserved.

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Copyright (c) 2003-2008, Sparta, Inc

All rights reserved.

Copyright (c) 2004, Cisco, Inc and Information Network Center of Beijing University of Posts and Telecommunications.

All rights reserved.

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003

oss@fabasoft.com

Author: Bernhard Penz [bernhard.penz@fabasoft.com]

Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) -----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY

THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) ----

Copyright (c) 2003-2008, Sparta, Inc

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Sparta, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 6: Cisco/BUPTNIC copyright notice (BSD) -----

Copyright (c) 2004, Cisco, Inc and Information Network Center of Beijing University of Posts and Telecommunications.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) ----

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003

oss@fabasoft.com

Author: Bernhard Penz bernhard.penz@fabasoft.com

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

RegExp

Copyright (c) 1986 by University of Toronto. Written by Henry Spencer. Not derived from licensed software.

RegExp License

```
// In case this isn't obvious from the later comments this is an ALTERED
// version of the software. If you like my changes then cool, but nearly
// all of the functionality here is derived from Henry Spencer's original
// work.

//
// This code should work correctly under both _SBCS and _UNICODE, I did
// start working on making it work with _MBCS but gave up after a while
// since I don't need this particular port and it's not going to be as
// straight forward as the other two.

//
// used everywhere. Certainly it's doable, but it's a pain.

// What's worse is that the current code will compile and run under _MBCS,
// only breaking when it gets wide characters thrown against it.

//
// I've marked at least one bit of code with #pragma messages, I may not
// get all of them, but they should be a start

//
// Guy Gascoigne - Piggford (ggp@bigfoot.com) Friday, February 27, 1998
// regcomp and regexec -- regsub and reerror are elsewhere
// @(#)regexp.c 1.3 of 18 April 87

//
// Copyright (c) 1986 by University of Toronto.

// Written by Henry Spencer. Not derived from licensed software.

// Permission is granted to anyone to use this software for any
// purpose on any computer system, and to redistribute it freely,
// subject to the following restrictions:
```

```
//  
// 1. The author is not responsible for the consequences of use of  
// this software, no matter how awful, even if they arise  
// from defects in it.  
//  
// 2. The origin of this software must not be misrepresented, either  
// by explicit claim or by omission.  
//  
// 3. Altered versions must be plainly marked as such, and must not  
// be misrepresented as being the original software.  
// *** THIS IS AN ALTERED VERSION. It was altered by John Gilmore,  
// *** hptoad!gnu, on 27 Dec 1986, to add and for word-matching  
// *** as in BSD grep and ex.  
// *** THIS IS AN ALTERED VERSION. It was altered by John Gilmore,  
// *** hptoad!gnu, on 28 Dec 1986, to optimize characters quoted with \.  
// *** THIS IS AN ALTERED VERSION. It was altered by James A. Woods,  
// *** ames!jaw, on 19 June 1987, to quash a regcomp() redundancy.  
// *** THIS IS AN ALTERED VERSION. It was altered by Geoffrey Noer,  
// *** THIS IS AN ALTERED VERSION. It was altered by Guy Gascoigne - Piggford  
// *** guy@wydrune.com, on 15 March 1998, porting it to C++ and converting  
// *** it to be the engine for the CRegexp class  
//  
// Beware that some of this code is subtly aware of the way operator  
// precedence is structured in regular expressions. Serious changes in  
// regular-expression syntax might require a total rethink.
```

